Engineering Art Galleries

Pedro J. de Rezende[1], Cid C. de Souza[1], Stephan Friedrichs[2], Michael Hemmer[3], Alexander Kröller[3], and Davi C. Tozoni[1]

[1]University of Campinas, Institute of Computing, Campinas, Brazil,
E-Mail: {rezende,cid}@ic.unicamp.br, davi.tozoni@gmail.com
[2]Max Planck Institute for Informatics, Saarbrücken, Germany,
E-Mail: sfriedri@mpi-inf.mpg.de.
[3]TU Braunschweig, IBR, Algorithms Group, Braunschweig, Germany,
E-Mail: mhsaar@gmail.com, kroeller@perror.de

arXiv:1410.8720v3 [cs.CG] 15 Feb 2016

Abstract

The Art Gallery Problem is one of the most well-known problems in Computational Geometry, with a rich history in the study of algorithms, complexity, and variants. Recently there has been a surge in experimental work on the problem. In this survey, we describe this work, show the chronology of developments, and compare current algorithms, including two unpublished versions, in an exhaustive experiment. Furthermore, we show what core algorithmic ingredients have led to recent successes.

Keywords: Art Gallery Problem, Computational Geometry, Linear Programming, Experimental Algorithmics.

1 Introduction

The Art Gallery Problem (AGP) is one of the classic problems in Computational Geometry (CG). Originally it was posed forty years ago, as recalled by Ross Honsberger [37, p.104]:

> "At a conference in Stanford in August, 1973, Victor Klee asked the gifted young Czech mathematician Václav Chvátal (University of Montreal) whether he had considered a certain problem of guarding the paintings in an art gallery. The way the rooms in museums and galleries snake around with all kinds of alcoves and corners, it is not an easy job to keep an eye on every bit of wall space. The question is to determine the minimum number of guards that are necessary to survey the entire building."

It should be noted that a slightly different definition is used today, where not only the walls of the gallery have to be guarded, but also the interior (this is indeed a different problem, see Figure 1a). AGP has received enormous attention from the CG community, and today no CG textbook is complete without a treatment of it. We give an overview on the most relevant developments in Section 2, after introducing the problem more formally.

Besides theoretical interest, there are practical problems that turn out to be AGP. Some are of these are straightforward, such as guarding a shop with security cameras, or illuminating an environment with few lights. For another example, consider a commercial service providing indoors laser scanning: Given an architectural drawing of an environment, say, a factory building, a high-resolution scan needs to be obtained. For that matter, the company brings in a scanner, places it on a few carefully chosen positions, and scans the building. As scanning takes quite a while, often in the range of several hours per position, the company needs to keep

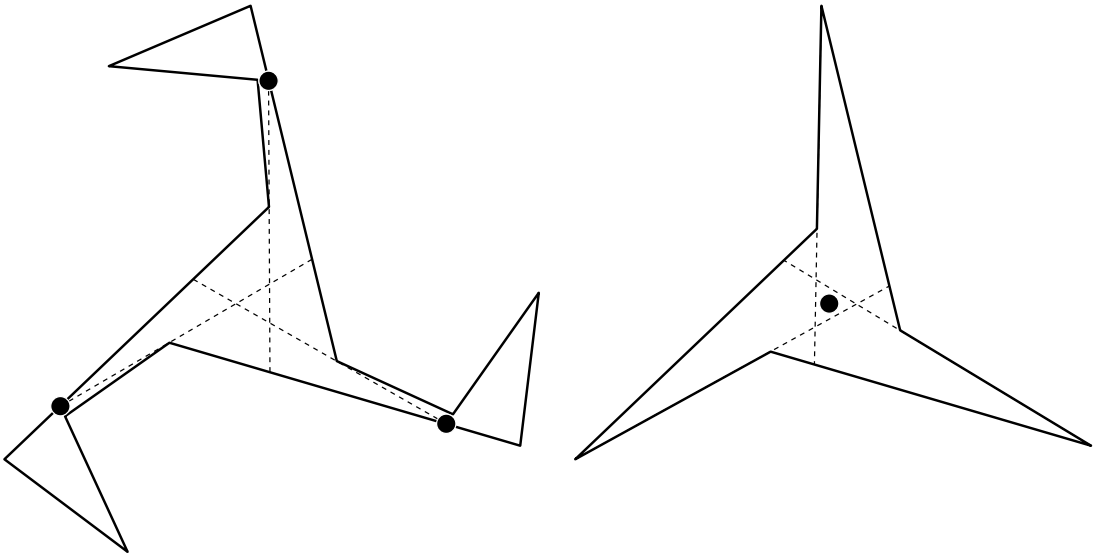

(a) Three guards suffice to cover the walls, but not the interior.

(b) One point guard covers the interior, but a vertex guard cannot

Figure 1: Edge cover and vertex guard variants have better and worse solutions than the classic AGP, respectively.

the number of scans as low as possible to stay competitive — this is exactly minimizing the number of guards (scan positions) that still survey (scan) the whole environment.

In this paper, we provide a thorough survey on experimental work in this area, i.e., algorithms that compute optimal or good solutions for AGP, including some problem variants. We only consider algorithms that have been implemented, and that underwent an experimental evaluation. During the past seven years, there have been tremendous improvements, from being able to solve instances with tens of vertices with simplification assumptions, to algorithm implementations that find optimal solutions for instances with several thousands of vertices, in reasonable time on standard PCs. We avoid quoting experimental results from the literature, which are difficult to compare to each other due to differences in benchmark instances, machines used, time limits, and reported statistics. Instead, we conducted a massive unified experiment with 900 problem instances with up to 5000 vertices, comparing six different implementations that were available to us. This allows us to pinpoint benefits and drawbacks of each implementation, and to exactly identify where the current barrier in problem complexity lies. Given that all benchmarks are made available, this allows future work to compare against the current state. Furthermore, for this paper, the two leading implementations were improved in a joint work between their respective authors, using what is better in each. The resulting implementation significantly outperforms any previous work, and constitutes the current frontier in solving AGP.

The remainder of this paper is organized as follows. In the next section, we formalize the problem and describe related work. In Section 3, we turn our attention to the sequence of experimental results that have been presented in the past few years, with an emphasis on the chronology of developments. This is followed by an experimental cross-comparison of these algorithms in Section 4, showing speedups over time, and the current frontier. In Section 5, we take an orthogonal approach and analyze common and unique ingredients of the algorithms, discussing which core ideas have been most successful. This is followed by a discussion on closely related problem variants and current trends in Section 6, and a conclusion in Section 7.

2 The Art Gallery Problem

Before discussing Art Gallery Problem (AGP) in detail, let us give a formal definition and introduce the necessary notation.

2.1 Problem and Definitions

We are given a polygon P, possibly with holes, in the plane with vertices V and $|V| = n$. P is *simple* if and only if its boundary, denoted by ∂P, is connected. For $p \in P$, $\mathcal{V}(p) \subseteq P$ denotes all points *seen* by p, referred to as the *visibility region* of p, i.e., all points $p' \in P$ that can be connected to p using the line segment $\overline{pp'} \subset P$. We call P *star-shaped* if and only if $P = \mathcal{V}(p)$ for some $p \in P$, the set of all such points p represents the *kernel* of P. For any $G \subseteq P$, we denote by $\mathcal{V}(G) = \bigcup_{g \in G} \mathcal{V}(g)$. A finite $G \subset P$ with $\mathcal{V}(G) = P$ is called a *guard set* of P; $g \in G$ is a guard. We say that g *covers* all points in $\mathcal{V}(g)$. The AGP asks for such a guard set of minimum cardinality.

Note that visibility is symmetric, i.e., $p \in \mathcal{V}(q) \iff q \in \mathcal{V}(p)$. The inverse of $\mathcal{V}(\cdot)$ describes all points that can see a given point p. This is easily confirmed to be

$$\mathcal{V}^{-1}(p) := \{q \in P : p \in \mathcal{V}(q)\} = \mathcal{V}(p) \ .$$

We use two terms to refer to points of P, making the discussion easier to follow. We call a point a *guard position* or *guard candidate* when we want to stress its role to be selected as part of a guard set. The second term comes from the fact that in a feasible solution, every point in $w \in P$ needs to be covered by some visibility polygon. We refer to such a point as *witness* when we use it as certificate for coverage.

Let $G, W \subseteq P$ be sets of guard candidates and witnesses such that $W \subseteq \mathcal{V}(G)$. The AGP variant were W has to be covered with a minimum number of guards, which may only be picked from G, can be formulated as an Integer Linear Program (ILP):

$$\begin{aligned}
\mathrm{AGP}(G, W) := \quad & \min \quad \sum_{g \in G} x_g && (1)\\
& \text{s.t.} \quad \sum_{g \in \mathcal{V}(w) \cap G} x_g \geq 1, \ \forall w \in W, && (2)\\
& \qquad x_g \in \{0, 1\}, \qquad \forall g \in G. && (3)
\end{aligned}$$

Essentially the model above casts the AGP variant in terms of a Set Covering Problem (SCP). But note that, depending on the choice of G and W, $\mathrm{AGP}(G, W)$ may have an infinite number of variables and/or constraints, i.e., be a semi- or doubly-infinite ILP. We discuss three major variants of AGP:

- The classic AGP definition, allowing for arbitrary *point guards*, i.e., allowing to place guards anywhere within P. It requires that all of P, boundary and interior, is guarded. This corresponds to $\mathrm{AGP}(P, P)$. We refer to this variant as "the" AGP.
- In $\mathrm{AGP}(V, P)$, all of P has to be guarded, but guards are restricted to be placed on vertices of P only. We refer to such guards as *vertex guards*. Trivially, a vertex guard solution is a solution for AGP as well, but the reverse is not necessarily true, see Figure 1b.
- The variant that Victor Klee actually described, i.e., where only the polygon's boundary needs to be guarded, is described by $\mathrm{AGP}(P, \partial P)$. A solution for $\mathrm{AGP}(P, P)$ also solves $\mathrm{AGP}(P, \partial P)$, but not vice versa (see Figure 1a).

There are many more AGP variants that deserve (and received) attention, however, these are the three versions that are mostly relevant for this paper.

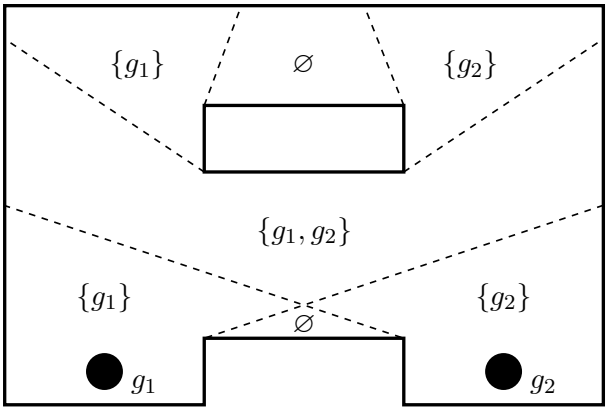

Figure 2: The visibility arrangement $\mathcal{A}(\{g_1, g_2\})$ induced by two guards g_1 and g_2 in a polygon with one hole.

In the following, unless explicitly stated otherwise, we use G and W to indicate discretized versions of the AGP. For example $\mathrm{AGP}(G, P)$ may refer to a (sub-)problem where all of P needs to be guarded, but a *finite* set of guard candidates is already known. Analogously, $\mathrm{AGP}(P, W)$ is the version where only a finite set of points needs to be covered, and $\mathrm{AGP}(G, W)$ is the fully discretized version.

The semi-infinite case $\mathrm{AGP}(G, P)$ provides some structure that can be exploited in algorithms. Consider Figure 2. We denote by $\mathcal{A}(G)$ the arrangement obtained by overlaying all visibility polygons $\mathcal{V}(g)$ for every $g \in G$. Every feature (face, edge, or vertex) of $\mathcal{A}(G)$ has a well-defined set of guards that completely sees it. Hence, any of those guards covers the entire feature, and we refer to them as Atomic Visibility Polygons (AVPs). We define a partial order on them as follows: For two faces f_1 and f_2, we define $f_1 \succ f_2$ if they are adjacent in $\mathcal{A}(G)$ and the set of guards seeing f_1 is a superset of those seeing f_2. The maximal (minimal) elements in the resulting poset are called light (shadow) AVPs. They can be exploited to solve the two semi-infinite cases: For given finite G, any subset of G that covers all shadow AVPs also covers P, hence is feasible for $\mathrm{AGP}(G, P)$. For finite W, there is always an optimal solution for $\mathrm{AGP}(P, W)$ that uses only guards in light AVPs of $\mathcal{A}(W)$, with at most one guard per AVP.

2.2 Related Work

Chvátal [13] was the first to prove the famous "Art Gallery Theorem", stating that $\lfloor n/3 \rfloor$ guards are sometimes necessary and always sufficient for polygons with n vertices. Later Fisk [32] came up with a simple proof for this theorem, beautiful enough to be included in the BOOK [1]. It also translates directly into a straightforward algorithm to compute such a set of guards. Note that for every n, there exist polygons that can be guarded by a single point (i.e., star-shaped polygons). So any algorithm producing $\lfloor n/3 \rfloor$ guards is merely a $\Theta(n)$-approximation. There are excellent surveys on theoretical results, especially those by O'Rourke [49] and Urrutia [61] should be mentioned.

Many variants of the problem have been studied in the past. For example, Kahn et al. [38] established a similar theorem using $\lfloor n/4 \rfloor$ guards for orthogonal polygons. There are variants where the characteristics of the guards have been changed. For example, *edge guards* are allowed to move along an edge and survey all points visible to some point on this edge. Instead of patrolling along an edge, *diagonal guards* move along diagonals, *mobile guards* are allowed to use both. See Shermer [56] for these definitions. Alternatively, variations on the guard's task have been considered, for example, Laurentini [44] required visibility coverage for the polygon's edges only. Another relevant problem related to the coverage of polygons considers watchman routes. A watchman route is a path in the interior of a polygon P such that every point of P

is seen by at least one point in the path. Therefore, a mobile guard moving along this path can do the surveillance of the entire polygon's area. Results on this problem can be found, for example, in Mitchell [47] and [51].

AGP and its variants are typically hard optimization problems. O'Rourke and Supowit [50] proved AGP to be NP-hard by a reduction from 3SAT, for guards restricted to be located on vertices and polygons with holes. Lee and Lin [45] showed NP-hardness also for simple polygons. This result was extended to point guards by Aggarwal [49]. Schuchardt and Hecker [55] gave NP-hardness proofs for rectilinear simple polygons, both for point and vertex guards. Eidenbenz et al. [25] established lower bounds on the achievable approximation ratio. They gave a lower bound of $\Omega(\log n)$ for polygons with holes. For vertex, edge and point guards in simple polygons, they established APX-hardness. For restricted versions, approximation algorithms have been presented. Efrat and Har-Peled [24] gave a randomized approximation algorithm with logarithmic approximation ratio for vertex guards. Ghosh [34] presented algorithms for vertex and edge guards only, with an approximation ratio of $O(\log n)$. For point guards Nilsson [48] gave $O(\mathrm{OPT}^2)$-approximation algorithms for monotone and simple rectilinear polygons. Also for point guards, Deshpande et al. [23] proposed one of the few existing approximation algorithms which is not constrained to a few polygon classes. See Ghosh [34] for an overview of approximation algorithms for the AGP. The first known exact algorithm for point guard problem was proposed by Efrat and Har-Peled [24] and has complexity $O((nc)^{3(2c+1)})$, where c is the size of the optimal solution. No experimental results with this algorithm have been reported so far. The exponential grow of the running time with c probably makes it useless to solve large non-trivial instances.

3 Timeline

After receiving mainly a theoretical treatment for over thirty years, several groups have started working on solving the AGP using the Algorithm Engineering methodology, aiming at providing efficient implementations to obtain optimal, or near-optimal, solutions.

Especially two groups, the Institute of Computing at the University of Campinas, Brazil, and the Algorithms Group at TU Braunschweig, Germany, developed a series of algorithms that substantially improve in what kind of instances can be solved efficiently. In this section, we give a chronological overview on these efforts, and describe the algorithms that were developed. It should be noted that all these approaches follow similar core ingredients, e.g., the AGP is treated as an infinite Set Covering Problem (SCP). As finite SCP instances can be solved reasonably fast in practice, the AGP is reduced to finite sets, and different techniques are employed to connect the finite and infinite cases.

3.1 Stony Brook 2007: AGP(P, P) heuristics, tens of vertices

Amit et al. [2] were among the first to experiment with a solver for AGP(P, P), see the journal version [3] and the PhD thesis by Packer [52] for extended presentations.

In this work, greedy algorithms are considered, following the same setup: A large set G of guard candidates is constructed, with the property that P can be guarded using G. Algorithms pick guards one after the other from G, using a priority function μ, until P is fully guarded. Both G and μ are heuristic in nature. The authors present 13 different strategies (i.e., choices for G and μ), and identify the three that are the best: In A_1, G consists of the polygon vertices, and of one additional point in every face of the arrangement obtained by adding edge extensions to the polygon. Priority is given to guards that can see the most of the currently unguarded other positions in G. The second strategy, A_2 follows the same idea. Additionally, after selecting a guard g, it adds $\mathcal{V}(g)$ to the arrangements and creates additional candidate

positions in the newly created faces. Finally, A_{13} employs a weight function ω on G, used as a random distribution. In each step, a point from G is selected following ω. Then, a random uncovered point p is generated, and all guard candidates seeing p get their weight doubled.

To produce lower bounds, greedy heuristics for independent witnesses (i.e., witnesses whose visibility regions do not overlap) are considered. Using a pool of witness candidates, consisting of the polygon's convex vertices and points on reflex-reflex edges, a witness set is constructed iteratively. In every step, the witness seeing the fewest other witness candidates is added, and dependent candidates are removed.

The authors conducted experiments with 40 input sets, including randomly generated as well as hand-crafted instances, with up to 100 vertices. Both simple polygons and ones with holes are considered. By comparing upper and lower bounds, it was found that the three algorithms mentioned above always produced solutions that are at most a factor 2 from the optimum. Algorithm A_1 was most successful in finding optimal solutions, which happened in 12 out of 37 reported cases.

3.2 Campinas 2007: AGP(V, P) for orthogonal simple polygons, hundreds of vertices

In 2007, Couto et al. [20, 21] focused on the development of an exact algorithm for the AGP with vertex guards, AGP(V, P), restricted to orthogonal polygons without holes. To the best of our knowledge, these works were the first in the literature to report extensive experimentation with an exact algorithm for a variant of the AGP. Early attempts to tackle the orthogonal AGP(V, P) also involved reductions to the SCP [27, 28] and aimed either to obtain heuristic solutions or to solve it exactly [57, 58]. However, experiments in these works only considered a few instances of limited sizes. In contrast, in the work of Couto et al., thousands of instances, some of which with 1000 vertices, were tested and later assembled into a benchmark, made publicly available for future comparisons [16], containing new classes of polygons including some very hard problem instances.

Moreover, in [21], the group in Campinas derived theoretical results that were later extended and proved to be instrumental to obtain exact solutions for more general variants of the AGP. They showed that AGP(V, P) can be solved through a single instance of the SCP by replacing the infinite set of points P by a finite set of suitably chosen witnesses from P.

The basic idea of the algorithm is to select a discrete set of points W in P and then solve the AGP variant whose objective consists in finding the minimum number of vertices sufficient to cover all points in W. This discretized AGP is then reduced to an SCP instance and modeled as an ILP. The resulting formulation is subsequently solved using an ILP solver, in their case, XPRESS. If the solution to the discretized version covers the whole polygon, then an optimal solution has been found. Otherwise, additional points are added to W and the procedure is iterated. The authors prove that the algorithm converges in a polynomial number of iterations, $O(n^3)$ in the worst case.

An important step of this exact algorithm is to decide how to construct the set of witnesses W. Couto et al. study various alternatives and investigated the impact on the performance of the algorithm. In the first version [20], a single method for selecting the initial discretization is considered, which is based on the creation of a regular grid in the interior of P. In the journal version [21], four new discretizations are proposed: *Induced grid* (obtained by extending the lines of support of the edges of the polygon), *just vertices* (comprised of all vertices of P), *complete AVP* (consisting of exactly one point in the interior of each AVP), and *reduced AVP* (formed by one point from each shadow AVP). The authors prove that, with the *shadow AVP* discretization, that it takes the algorithm only one iteration to converge to an optimal solution for the orthogonal AGP.

The first experimental results were largely surpassed by those reported in the journal version. Besides introducing new discretizations, the *shadow AVP* discretization increased the polygon

sizes fivefold (to 1000 vertices). In total, almost 2000 orthogonal polygons were tested, including von Koch polygons, which give rise to high density visibility arrangements and, as a consequence, to larger and harder to solve SCP instances.

The authors highlight that, despite the fact that the visibility polygons and the remaining geometric operations executed by the algorithm can be computed in polynomial time, in practice, the preprocessing phase (i.e. geometric operations such as visibility polygon computation) is responsible for the majority of the running time. At first glance, this is surprising since the SCP is known to be NP-hard and one instance of this problem has to be solved at each iteration. However, many SCP instances are easily handled by modern ILP solvers, as is the case for those arising from the AGP. Furthermore, the authors also observe that, when reasonable initial discretizations of the polygon are used, the number of iterations of the algorithm is actually quite small.

Knowing that the *reduced AVP* discretization requires a single iteration, albeit an expensive one timewise, the authors remark that a trade-off between the number of iterations and the hardness of the SCP instances handled by the ILP solver should to be sought. Extensive tests lead to the conclusion that the fastest results were achieved using the *just vertices* discretization since, although many more iterations may be required, the SCP instances are quite small.

3.3 Torino 2008: AGP$(P, \partial P)$, hundreds of vertices

In 2008, Bottino and Laurentini [9] proposed a new algorithm for the AGP variant whose objective consists in only covering the edges of a polygon P, denoted AGP$(P, \partial P)$. Hence, in this version, coverage of the interior of P is not required. Despite being less constrained than the original AGP, the AGP$(P, \partial P)$ was proven to be NP-hard [44]. In this context, the authors presented an algorithm capable of optimally solving AGP$(P, \partial P)$ for polygons with and without holes provided the method converges in a finite number of steps. This represents a significant improvement in the search for optimal solutions for the AGP.

The algorithm by Bottino and Laurentini works iteratively. First, a lower bound specific for P is computed. The second step consists of solving an instance of the so called Integer Edge Covering Problem (IEC). In this problem, the objective is also to cover the whole boundary of the polygon with one additional restriction: each edge must be seen entirely by at least one of the selected guards. It is easy to see that a solution to the IEC is also viable for AGP$(P, \partial P)$ and, consequently, its cardinality is an upper bound for the latter. After obtaining a viable solution, the gap between the upper and lower bounds is checked. If it is zero (or less than a predefined threshold) the execution is halted. Otherwise, a method is used to find *indivisible* edges, which are edges that are entirely observed by one guard in some or all optimal solutions of AGP$(P, \partial P)$. The identification of these edges can be done in polynomial time from the visibility arrangement. After identifying them, those classified as not indivisible are split and the process starts over.

Tests were performed on approximately 400 random polygons with up to 200 vertices. The instances were divided into four classes: simple, orthogonal, random polygons with holes and random orthogonal polygons with holes. Reasonable optimality percentages were obtained using the method. For instance, on random polygons with holes, optimal results were achieved for 65% of the instances with 60 vertices. In cases where the program did not reach an optimal solution (due to the optimality gap threshold or to timeout limits), the final upper bound was, on average, very close to the lower bound computed by the algorithm. On average, for all classes of polygons, the upper bound exceeded the lower bound by $\sim 7\%$.

3.4 Campinas 2009: AGP(V, P) for simple polygons, thousands of vertices

Couto et al. went on to study how to increase the efficiency of the algorithm for AGP(V, P) proposed in the works discussed in Section 3.2. A complete description of their findings can be found in a 2011 paper [19], with a preliminary version available as a 2009 technical report [15]. The basic steps of the algorithm are explained in [17], and are illustrated in the companion video [18].

Compared to the previous works by the same authors, the new algorithm was extended to cope with more general classes of polygons, still without holes, but now including non-orthogonal polygons. Experiments on thousands of instances confirmed the robustness of the algorithm. A massive amount of data was subsequently made publicly available containing the entire benchmark used for these tests, see also Section 4.1.

Essentially, some implemented procedures were improved relative to the approach in [21] to enable handling non-orthogonal polygons. Moreover, two new initial discretization techniques were considered. The first one, called *single vertex*, consists in the extreme case where just one vertex of the polygon forms the initial discretized set W. As the second strategy, named *convex vertices*, W comprises all convex vertices of P.

The authors made a thorough analysis of the trade-off between the number and nature of the alternative discretization methods and the number of iterations. Their tests were run on a huge benchmark set of more than ten thousand polygons with up to 2500 vertices. The conclusion was that the decision over the best discretization strategy deeply depends on the polygon class being solved. As anticipated, the fraction of time spent in the preprocessing phase was confirmed to be large for sizable non-orthogonal polygons and even worse in the case of von Koch and random von Koch polygons. Moreover, while using shadow AVPs as the initial discretization produces convergence after just one iteration of the algorithm, the resulting discretization set can, in this case, be so large that the time cost of the preprocessing phase overshadows the solution of the ensued SCP. For this reason, the *just vertices* strategy lead to the most efficient version of the algorithm, in practice, as the small SCP instances created counterbalanced the larger number of iterations for many polygon classes.

3.5 Braunschweig 2010: Fractional solutions for AGP(P, P), hundreds of vertices

In 2010, Baumgartner et al. [7] (see Kröller et al. [41] for the journal version) presented an exact algorithm for the fractional variant of AGP(P, P). In it, solutions may contain guards $g \in P$ with a fractional value for x_g. This corresponds to solving a Linear Program (LP), namely the LP relaxation of AGP(P, P) which is obtained by replacing Constraint (3) of AGP(G, W) with

$$0 \leq x_g \leq 1 \quad \forall g \in G. \tag{4}$$

We denote by AGP$_{\mathbb{R}}(G, W)$ the LP relaxation of AGP(G, W). Note that this is the first exact algorithm where neither guard nor witness positions are restricted.

The authors present a primal-dual approach to solve the problem. They notice that the fractional AGP$_{\mathbb{R}}(G, W)$ can be easily solved using an LP solver, provided G and W are finite and not too large. The proposed algorithm picks small, carefully chosen sets for G and W. It then iteratively extends them using cutting planes and column generation:

Cutting Planes If there is an uncovered point $w \in P \setminus W$, this corresponds to a violated constraint of AGP$_{\mathbb{R}}(P, P)$, so w is added to W. Otherwise the current solution is feasible for AGP$_{\mathbb{R}}(G, P)$, and hence an upper bound of AGP$_{\mathbb{R}}(P, P)$. We also refer to this part as *primal separation*.

Column Generation A violated constraint of the dual of $\mathrm{AGP}_{\mathbb{R}}(P, P)$ corresponds to a guard candidate $g \in P \backslash G$ that improves the current solution, and g is added to G. Otherwise the current solution optimally guards the witnesses in W, i.e. is optimal for $\mathrm{AGP}_{\mathbb{R}}(P, W)$, and hence provides a lower bound for $\mathrm{AGP}_{\mathbb{R}}(P, P)$. We also refer to this part as *dual separation*.

It can be shown that, if the algorithm converges, it produces an optimal solution for $\mathrm{AGP}_{\mathbb{R}}(P, P)$. Furthermore, the authors use the algorithm for the integer AGP, but only insofar that LP solutions sometimes are integer by chance, but without any guarantee.

The algorithm has many heuristic ingredients, e.g., in the choice of initial G and W and the placement strategy for new guards and witnesses. The authors conducted an exhaustive experiment, with 150 problem instances with up to 500 vertices. They compared different strategies for all heuristic ingredients of the algorithm. There were four separation strategies: (1) Focusing on upper bounds by always running primal separation, but dual only when primal failed. (2) Focusing on lower bounds, by reversing the previous one. (3) Always running both separators, in the hope of quickly finishing. (4) Alternating between foci, by running primal separation until an upper bound is found, then switching to running dual separation until a lower bound is found, and repeating. There were four different separators, i.e., algorithms to select new candidates for G resp. W; these included selecting the point corresponding to a maximally violated constraint, selecting points in all AVPs, a greedy strategy to find independent rows (columns) with a large support, and placing witnesses on AVP edges. For initial choice of G and W, four heuristics were used: (1) Using all polygon vertices, (2) starting with an empty set (for implementation reasons, a single point had to be used here), (3) selecting half the vertices to keep the set smaller but still allowing full coverage, and finally two strategies based on the work by Chwa et al. [14]. Here, G is initialized to use all reflex vertices, and W is initialized to have a witness on every polygon edge that is incident to a reflex vertex.

The trial consisted of over 18,000 runs, allowing for a direct comparison of individual parameter choices. It was found that many instances could be solved optimally within 20 minutes. This happened for 60% of the 500-vertex polygons, and 85% of the 100-vertex polygons. Other findings included the importance of the initial solution, where the best strategy (the Chwa-inspired one) led to an overall speedup factor of 2. The best primal and dual separators were identified in a similar way. Furthermore, the authors were first to observe the bathtub-shaped runtime distribution that is still prominent in today's algorithms: Either the algorithm finishes very quickly, usually in the first few seconds, with an optimal solutions, or it engages in an excruciatingly slow process to find good guards and witnesses, often failing to finish within time.

3.6 Torino 2011: $\mathrm{AGP}(P, P)$, tens of vertices

In 2011, Bottino et al. [10] improved their previous work (see Section 3.3) by applying similar ideas to solve the original AGP rather than the $\mathrm{AGP}(P, \partial P)$. The objective was to develop an algorithm capable of finding nearly-optimal solutions for full polygon coverage, since, at that time, there was a lack of practical methods for this task.

The first step of the presented technique consists of using the algorithm discussed in Section 3.3, which allows for obtaining a lower bound and also multiple optimal solutions for the $\mathrm{AGP}(P, \partial P)$. These are then tested in the search for a coverage of the entire polygon. According to the authors, if such solution exists, it is automatically a nearly optimal one for $\mathrm{AGP}(P, P)$. If a viable solution is not among those, guards are then added using a greedy strategy until a feasible solution is found.

It should be noted that there are worst-case instances for $\mathrm{AGP}(P, P)$ that only possess a single optimal solution, where no characterization of the guard positions is known. Therefore

this algorithm, and none of the subsequent ones presented in this paper, can guarantee to find optimal solutions. This common issue is discussed in more detail in Section 6.2.

For the experiments presented in [10], 400 polygons with sizes ranging from 30 to 60 vertices were examined. As in the previous work, the following classes were tested: simple, orthogonal, random polygons with holes and also random orthogonal polygons with holes. Guaranteed optimal solutions were found in about 68% of the polygons tested. Note that, in about 96% of the cases, the solution found for $\mathrm{AGP}(P, \partial P)$ in the first step of this algorithm was also viable for $\mathrm{AGP}(P, P)$. In addition, the authors also implemented the most promising techniques by Amit et al. (see Section 3.1), in order to enable comparison between both works. As a result, this technique was more successful than the method by Amit et al. considering the random polygons tested.

3.7 Braunschweig 2012: $\mathrm{AGP}(P, P)$, hundreds of vertices

In 2012, the primal-dual method introduced by the Braunschweig group was extended. The previous version could find optimal point guards, but only for the LP relaxation which allows fractional guards. Integer solutions could only be found by chance. Now, two ingredients were added to find integer solutions: An ILP-based routine and cutting planes. See Friedrichs [33] for a detailed discussion on the cutting planes, and Fekete et al. [30, 31] for the combined approach. As it turned out, this algorithm could solve the classic problem $\mathrm{AGP}(P, P)$ on instances of several hundreds of vertices with holes, a factor 10 more than in previous work.

The 2012 algorithm switches between primal and dual phases. In the primal phase, feasible solutions are sought, i.e., upper bounds. Unlike the 2010 version, now only integer solutions are considered. For the current set G of guards and W of witnesses, $\mathrm{AGP}(G, W)$ is solved optimally using an ILP formulation. The visibility overlay $\mathcal{A}(G)$ is scanned for insufficiently covered spots, and additional witnesses are generated accordingly. The primal phase ends when no new witnesses are generated, with a feasible integer solution for $\mathrm{AGP}(G, P)$, and hence an upper bound for $\mathrm{AGP}(P, P)$. In the dual phase, new guard positions are found using the dual arrangement $\mathcal{A}(W)$. For that, a dual solution is needed, which is provided by solving the LP relaxation $\mathrm{AGP}_{\mathbb{R}}(G, W)$. The dual phase ends with an optimal solution for $\mathrm{AGP}_{\mathbb{R}}(P, W)$, which is a lower bound for $\mathrm{AGP}_{\mathbb{R}}(P, P)$, and hence also $\mathrm{AGP}(P, P)$. The procedure computes a narrowing sequence of upper bounds for $\mathrm{AGP}(P, P)$ and lower bounds for $\mathrm{AGP}_{\mathbb{R}}(P, P)$, leaving the issue of closing the integrality gap between them. This may lead to terminating with a suboptimal solution, however with a provided lower bound. As a leverage against this shortcoming, cutting planes are employed to raise the lower bounds [33]. Two classes of facet-defining inequalities for the convex hull of all feasible integer solution of $\mathrm{AGP}(G, W)$ are identified. While the NP-hardness of AGP indicates that it is hopeless to find a complete polynomial-size facet description, it is shown that the new inequalities contain a large set of facets, including all with coefficients in $\{0, 1, 2\}$, see also [6]. The dual phase is enhanced with separation routines for the two classes, consequently improving the lower bounds, and often allowing the algorithm to terminate with provably optimal solutions.

To evaluate this work, experiments were conducted on four different polygon classes, sized between 60 and 1000 vertices. These included both orthogonal and non-orthogonal instances, both with and without holes, and polygons where optimal solutions cannot use vertex guards. Different parametrizations of the algorithms were tested, and it was found that the ILP-based algorithm itself (without applying cutting planes) could identify good integer solutions, sometimes even optimal ones, and considerably surpassed the previous 2010 version. The algorithm was able to find optimal solutions for 500-vertex instances quite often. Instances with 1000 vertices were out of reach though.

3.8 Campinas 2013: AGP(P, P), hundreds of vertices

The work by Tozoni et al. [60] generalizes to AGP(P, P) the ideas developed for AGP(V, P) by Couto et al. (see Section 3.4). The paper proposes an algorithm that iteratively generates upper and lower bounds while seeking to reach an exact solution. Extensive experiments were carried out which comprised 1440 simple polygons with up to 1000 vertices from several classes, all of which were solved to optimality in a matter of minutes on a standard desktop computer. Up to that point in time, this was the most robust and effective algorithm available for AGP(P, P), for simple polygons. The restriction to simple polygons in this version as well as earlier versions of the Campinas branch originates from the fact that no visibility algorithm for general polygons was available to the group in Campinas, yet.

The algorithm generates, through a number of iterations, lower and upper bounds for the AGP(P, P) through the resolution of the two semi-infinite discretized variants of the original AGP, namely AGP(P, W) (asking for the minimum number of guards that are sufficient to cover the finite set W of witnesses) and AGP(G, P) (computing the minimum number of guards from G that are sufficient to cover P). Notice that in these variants, either the witness or the guard candidate set is infinite, preventing the formulation of these problem variants as an ILP. However, remarkable results [60] show that both variants can be reduced to a compact set covering problem.

To solve AGP(P, W) instance, the algorithm constructs $\mathcal{A}(W)$, and chooses the vertices of the light AVPs to become part of the guard candidates set G. Assuming that $|W|$ is bounded by a polynomial in n, the same holds for $|G|$. Therefore, the SCP instance corresponding to AGP(G, W) admits a compact ILP model. Tozoni et al. showed that an optimal solution for AGP(P, W) can be obtained by solving AGP(G, W). Thus, the algorithm computes a lower bound for AGP(P, P) using an ILP solver.

Now, to produce an upper bound for AGP(P, P), an idea similar to the one developed by Couto et al. [19] to solve the AGP(V, P) is used. The procedure starts with the same sets G and W used for the lower bound computation. The AGP(G, W) is solved as before. If the optimal solution found in this way covers P, then it is also feasible for AGP(P, P) and provides an upper bound for the problem. Otherwise, new witnesses are added to the set W and the procedure iterates. The upper bound procedure is known to converge in a number of iterations that is polynomial in n.

The lower and upper bound procedures are repeated until the gap between the two bounds reaches zero or a predefined time limit is reached. For certain initial discretization sets and strategies for updating the witness set, one can construct fairly simple instances that lead the algorithm to run indefinitely. Therefore, it remains an important open question whether there exists a discretization scheme that guarantees that the algorithm always converges, see also Section 6.2.

An important step of this algorithm, which greatly affects the performance of the final program, is how the initial witness set should be chosen and updated throughout the iterations. Two initial discretizations were tested in practice and are worth noting. The first one, called *Chwa-Points*, is based on the work by Chwa et al. [14] and chooses the middle points of reflex-reflex edges and the convex vertices that are adjacent to reflex vertices. This is similar to the initialization used in [7, 41]. The second, called *Convex-Vertices*, comprises all convex vertices of P.

The computational results obtained by this algorithm confirmed its robustness. A total of 1440 instances were tested from different polygon classes, including simple, orthogonal and von Koch ones. Optimal solutions were found for all of them. Also, comparisons with previous published papers showed that the algorithm was effective and far more robust than its competitors. Experiments with different initial witness sets revealed that, on average, *Chwa Points* attained the best results. However, on von Koch Polygons, *Convex Vertices* performed better.

3.9 Campinas 2013 (Journal version): AGP(P, P), thousands of vertices

After presenting an algorithm for AGP with point guards in spring 2013 (see Section 3.8), the research group from Campinas continued working on the subject. In this context, improvements were implemented, including the development of their own visibility algorithm that was also able to handle polygons with holes, giving rise to a new version of the algorithm [59]. The resulting implementation is able to solve polygons with thousands of vertices in a few minutes on a standard computer.

Several major improvements were introduced in order to reach this redesigned version. Among them, a Lagrangian Heuristic method (see Section 5.2.1) was implemented to help the ILP solver expedite the computation of optimal solutions for SCP instances. Moreover, a procedure for removing redundant variables and constraints from the SCP formulation was also used to speed up the ILP resolution process.

One of the most effective changes consisted in reversing the point of view of visibility testing from the perspective of the guards to that of the witnesses. Since these are fewer in number and their arrangement has already been computed, much of the geometric computation is simplified.

In the end, 2440 instances were tested and optimal solutions were found for more than 98% of them. The test bench included several different classes of polygons, with and without holes, with up to 2500 vertices. Besides the classes tested in the previous version [60], the authors also used a newly created benchmark instance for polygons with holes (see also Section 4.1) and the spike polygons presented by Kröller et al. [41]. Also, comparisons were made with the work of Kröller et al. and an analysis of the effects obtained from different discretizations for the initial witness set were presented. Moreover, the authors evaluated the impact of using a Lagrangian heuristic on the overall performance of the method and concluded that it reduces the average execution time in most of the cases.

3.10 Braunschweig 2013 (Current version): AGP(P, P), thousands of vertices

A deeper runtime analysis of the former code from 2012 revealed that the main bottlenecks where the geometric subroutines, namely (i) the computation of visibility polygons (an implementation of a $O(n \log n)$ rotational sweep as in Asano [4]), (ii) the overlays of these visibility polygons to form $\mathcal{A}(G)$ and $\mathcal{A}(W)$ ($O(n^2 m^2 \log(nm))$, where m is the size of G resp. W), and (iii) point location algorithms to determine membership in AVPs. This was somewhat surprising as all of these algorithms have fairly low complexity, especially when compared to LP solving (worst-case exponential time when using the Simplex algorithm) and ILP solving (NP-hard in general, and used to solved the NP-hard Set Cover problem). Still the geometric routines made up for over 90% of the runtime.

The group in Braunschweig focused on the improvement of these geometric subroutines: (i) A new Computational Geometry Algorithms Library [12] (CGAL) package for visibility polygon computation was developed in Braunschweig [36], which contains a new algorithm named triangular expansion [11]. Though the algorithm only guarantees an $O(n^2)$ time complexity, it usually performs several magnitudes faster than the rotational sweep. (ii) The code now uses the lazy-exact kernel [54], which delays (or even avoids) the construction of exact coordinates of intersection points as much as possible. The impact is most evident in the construction of the overlays, which contain many intersection points. (iii) The algorithm was restructured to allow a batched point location [62, Sec. 3][1] of all already existing guards (or witnesses) with respect to a new visibility polygon at once.

[1] This is an an $O((n + m) \log n)$ sweep line algorithm, where n is the number of polygon vertices and m the number of query points.

The new code now runs substantially faster, allowing it to solve much larger instances than the previous one. This paper contains the first experimental evaluation of this new algorithm. Section 4 contains results from running the algorithm and comparing it to the other approaches presented here. Section 5 discusses the speedup obtained by the new subroutines.

3.11 Campinas and Braunschweig 2013 (Current version): AGP(P, P), thousands of vertices

This implementation is the result of a joint effort by the Braunschweig and the Campinas groups. With the intent of achieving robustness, its core is the algorithm from Campinas (Section 3.9), refitted with optimizations from Braunschweig that greatly improved its efficiency.

The new code now also uses the lazy exact kernel (cf. 3.10) of CGAL and the triangular expansion algorithm [11] of the new visibility package [36] of CGAL. While the impact of the new visibility polygon algorithm was huge for both approaches the usage of the lazy kernel was also significant since the overlays in the approach of Campinas contain significantly more intersection points. To see more about how changes in kernel and visibility affect the solver, consult Section 5.

Moreover, the current version of Campinas also includes new approaches on the algorithm side. One of the ideas developed was to postpone the computation of an upper bound (solving AGP(G, P)) to the time that a good lower bound, and, consequently, a "good" set of guard candidates is obtained. This can be done by repeatedly solving only AGP(P, W) instances until an iteration where the lower bound is not improved is reached. This situation possibly means that the value obtained will not change much in the next iterations. It also increases the chances that the first viable solution found is also provably optimal, which automatically reduces the number of AGP(G, P) instances which must be re-solved.

Other changes that are the inclusion of a new strategy for guard positioning, where only one interior point from each light AVP is chosen to be part of the guard candidate set (instead of all its vertices), and the possibility of using IBM ILOG CPLEX Optimization Studio [22] (CPLEX) solver instead of XPRESS.

This new version was tested in experiments conducted for this paper, using 900 problem instances ranging from 200 to 5000 vertices. Section 4 presents the obtained results in detail. The implementation proved to be efficient and robust for all classes of polygons experimented.

4 Experimental Evaluation

To assess how well the AGP can be solved using current algorithms, and how their efficiency has developed over the last years, we have run exhaustive experiments. The experiments involve all algorithms for which we could get working implementations, and were conducted on the same set of instances and on the same machines, see Section 4.2. We refrain from providing comparisons based on numbers from the literature.

We had several snapshots from the Braunschweig and Campinas code available, these are:

- For Braunschweig, the versions from 2010 (Section 3.5), 2012 (Section 3.7), and 2013 (3.10). These will be referred to as `BS-2010`, `BS-2012`, and `BS-2013`, respectively.
- For Campinas, the version from 2009 (Section 3.4), and the two snapshots from 2013 (Sections 3.8 and 3.9). These will be referred to as `C-2009`, `C-2013.1` and `C-2013.2`, respectively.
- The latest version is the combined approach from Campinas and Braunschweig that was obtained during a visit of Davi C. Tozoni to Braunschweig (Section 3.11), which we refer to as `C+BS-2013`.

The older versions have already been published, for these we provide a unified evaluation. The versions `BS-2013` and `C+BS-2013` are, as of yet, unpublished.

4.1 AGPLib

For the performed experiments, several classes of polygons were considered. The majority of them were collected from AGPLib [16], which is a library of sample instances for the AGP, consisting of various classes of polygons of multiple sizes. They include the test sets from many previously published papers [20, 21, 19, 60, 59, 7, 41].

To find out more about how each of the classes was generated, see [20] and [59]. Below, we show a short description of the six classes of instances considered in this survey; all of them are randomly generated:

"simple" Random non-orthogonal simple polygons as in Figure 3a.

"simple-simple" Random non-orthogonal polygons as in Figure 3b. These are generated like the "simple" polygon class, but with holes. The holes are also generated like the first class and randomly scaled and placed until they are in the interior of the initial polygon.

"ortho" Random floorplan-like simple polygons with orthogonal edges as in Figure 3c.

"ortho-ortho" Random floorplan-like orthogonal polygons as in Figure 3d. As the simple-simple class, these polygons are generated by using one polygon of the ortho class as main polygon, and then randomly scaling and translating smaller ortho polygons until they are contained within the main polygon's interior, where they are used as holes.

"von Koch" Random polygons inspired by randomly pruned Koch curves, see Figure 3e.

"spike" Random polygons with holes as in Figure 3f. Note that this class is specifically designed to provide polygons that encourage placing point guards in the intersection centers. It has been published along with the `BS-2010` algorithm (Section 3.5), which was the first capable of placing point guards.

4.2 Experimental Setup

The experiments were run on identical PCs with eight-core Intel Core i7-3770 CPUs at 3.4 GHz, 8 MB cache, and 16 GB main memory running a 64-bit Linux 3.8.0 kernel. All algorithms used version 4.0 of Computational Geometry Algorithms Library [12] (CGAL) and IBM ILOG CPLEX Optimization Studio [22] (CPLEX) 12.5. The only component using concurrency is the ILP solver CPLEX, everything else was single-threaded. For each polygon-class/complexity combination, we tested 30 different polygons. Each test run had a runtime limit of 20 minutes.

4.3 Results

Historically, the two lines of algorithms have been working towards AGP from different angles. Campinas focused on binary solutions, which initially came at the expense of being limited to given guard discretization, like vertex guards: $\mathrm{AGP}(V, P)$. The Braunschweig work started with point guards, but the price were fractional solutions: $\mathrm{AGP}_{\mathbb{R}}(P, P)$. It only became possible to reliably solve the binary AGP with point guards, $\mathrm{AGP}(P, P)$, with the `BS-2012` and `C-2013.1` algorithms.

Therefore, we first sketch progress for the AGP with vertex guards and the fractional AGP before discussing the experimental results for AGP with point guards itself.

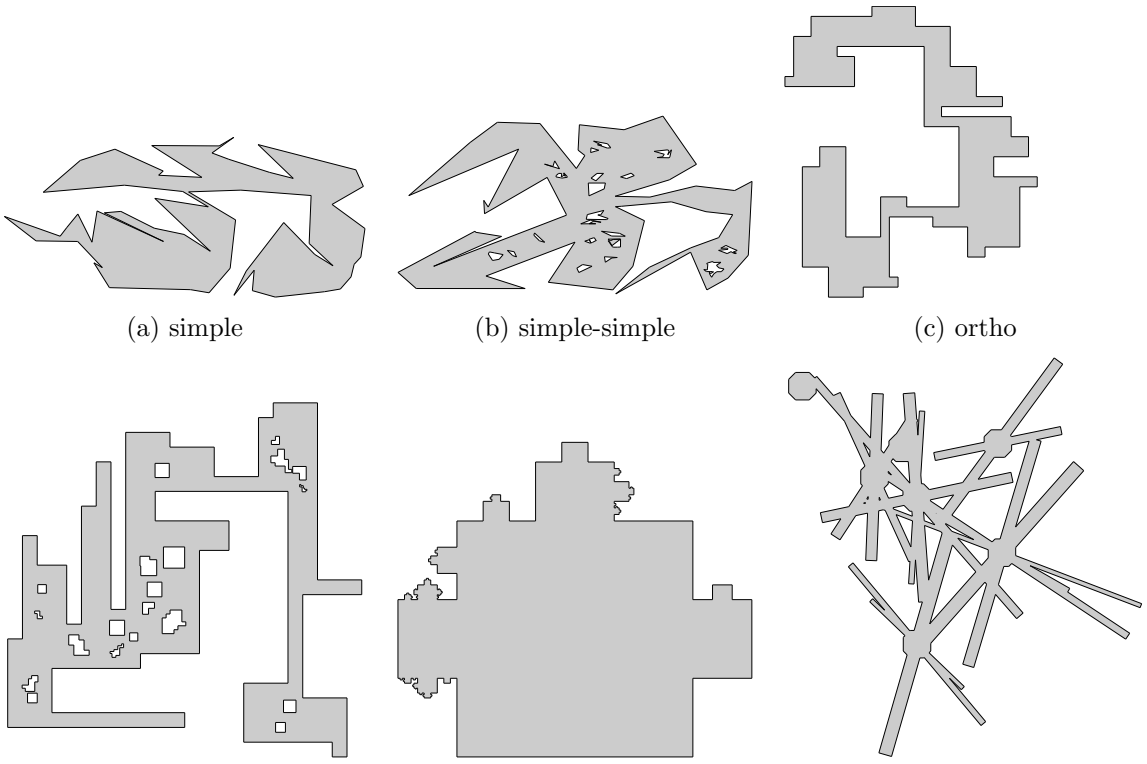

(a) simple (b) simple-simple (c) ortho

(d) ortho-ortho (e) vonKoch (f) spike

Figure 3: Example instances of different polygon classes.

4.3.1 Vertex Guards

AGP(V, P) is one of the semi-infinite variants. We believe it to be a considerably simpler problem than AGP(P, P) for two reasons: (1) Both variants are NP-hard, but we know AGP(V, P) is in NP, which is uncertain for AGP(P, P) as it is unknown if there is a polynomial-size representation of guard locations. (2) Experience and experimental results indicate that finding good guard candidates is the hardest part of the problem and leads to many iterations; but for AGP(V, P) we only have to place witnesses and solve SCP instances, which is usually possible in a comparably short time frame with a good ILP solver. The first experimental work on AGP(V, P) was Campinas 2007, but unfortunately, the implementation is no longer available.

Table 1 shows optimality rates, i.e., how many of the instances each implementation could solve, given a 20 minute time limit per instance. The polygons were grouped in two categories: those without holes, including the instances classes **simple**, **ortho** and **von Koch**, and those with holes composed by the instances in the classes **simple-simple**, **ortho-ortho** and **spikes**.

	polygons without holes					**polygons with holes**				
	200	500	1000	2000	5000	200	500	1000	2000	5000
`C-2009`	100.0	100.0	100.0	100.0	63.3	–	–	–	–	–
`C-2013.1`	100.0	100.0	100.0	100.0	63.3	–	–	–	–	–
`BS-2013`	100.0	100.0	100.0	100.0	88.9	100.0	97.8	67.8	77.8	66.7
`C-2013.2`	100.0	100.0	100.0	100.0	37.8	100.0	100.0	66.7	65.6	0.0
`C+BS-2013`	100.0	100.0	100.0	100.0	100.0	100.0	100.0	77.8	88.9	66.7
`BS-2010` *	100.0	100.0	100.0	100.0	33.3	100.0	100.0	100.0	98.9	0.0

Table 1: Optimality rates for vertex guards. Notice that `BS-2010` finds fractional vertex guard solutions, whereas the others find integer ones.

Class	n	Speedup Factor					
		C-2009	BS-2010	C-2013.1	C-2013.2	BS-2013	C+BS-2013
Simple	200	1.00	0.66	1.03	1.21	7.54	**6.75**
	500	1.00	0.66	1.01	1.02	7.79	**10.21**
	1000	1.00	0.66	1.02	0.95	8.03	**14.65**
	2000	1.00	0.68	1.00	0.90	10.24	**18.97**
	5000	–	–	–	–	–	–
Orthogonal	200	1.00	0.64	1.01	1.05	**6.46**	6.15
	500	1.00	0.63	1.01	0.98	6.67	**10.82**
	1000	1.00	0.65	1.00	0.92	7.75	**15.67**
	2000	1.00	0.65	0.98	0.82	9.57	**19.52**
	5000	1.00	0.72	1.00	0.75	12.63	**28.64**
von Koch	200	1.00	0.38	1.02	1.33	2.09	**3.45**
	500	1.00	0.44	1.09	1.37	1.86	**4.27**
	1000	1.00	0.60	0.95	1.40	1.95	**4.75**
	2000	1.00	0.92	1.34	1.39	2.67	**6.18**
	5000	–	–	–	–	–	–

Table 2: Speedup for vertex guards. Numbers indicate how many times faster than C-2009 later implementations became, computed as log-average. The comparison is only possible when there is at least one instance of the group that was solved by all considered solvers. This table is restricted to simple polygons, since C-2009 does not support polygons with holes.

The Campinas versions prior to C-2013.2 could not deal with holes in input polygons, so these entries are empty. It should also be noted that BS-2010 solves the easier case of fractional vertex guards. It is clearly visible how all algorithms (including the five-year-old C-2009) can solve all simple polygons with up to 2000 vertices as well as most simple 5000-vertex polygons. For instances with holes, however, the solution percentages of all algorithms (except BS-2010 which solves an easier problem) start significantly dropping at 1000 vertices. This demonstrates two effects: First, for smaller sizes, the problem is easier to solve as the search for good guard candidates is unnecessary. Second, for larger sizes, finding optimal solutions to large instances of the NP-hard SCP dominate, resulting in a computational barrier. The difficulty to handle large SCP instances also shows up when we consider the results of the Campinas codes C-2013.1 and C-2013.2. As the size of the polygons increases and the SCPs to be solved grow in complexity, the Lagrangian heuristic employed by C-2013.2 version uses more computational time but does not help the ILP solver to find optimal solutions for the AGP(G,W) instances, due to the deterioration of the primal bounds. This inefficiency causes a decrease in the solver's performance, as can be seen in the optimality rate shown in Table 1 for simple polygons with 5000 vertices. In this case, if C-2013.2 did not use the Lagrangian heuristic by default, a result at least similar to that obtained by C-2013.1 would be expected.

The high solution rates allow us to directly analyze the speedup achieved over time. Table 2 shows how much faster than C-2009 later algorithms could solve the problem. The shown numbers are log-averages over the speedup against C-2009 for all instances solved by both versions. As C-2009 cannot process holes, this analysis is restricted to simple polygons. It is clearly visible that BS-2013 is about five times faster then C-2009, and the changes from C-2013.2 to C+BS-2013 led to a speedup factor of about seven. These stem from a number of changes between versions, however, roughly a factor 5 can be attributed to improvements in geometric subroutines — faster visibility algorithms, lazy-exact CGAL kernel, reduced point constructions. We discuss the influence of geometry routines in Section 5.1.

		polygons without holes					polygons with holes				
		200	500	1000	2000	5000	200	500	1000	2000	5000
OPT	BS-2010	55.6	27.8	14.4	3.3	0.0	54.4	32.2	28.9	30.0	0.0
	BS-2012	53.3	30.0	11.1	4.4	0.0	54.4	32.2	27.8	31.1	0.0
	BS-2013	56.7	24.4	12.2	1.1	0.0	50.0	31.1	27.8	31.1	33.3
5% gap	BS-2010	93.3	100.0	100.0	100.0	33.3	96.7	100.0	98.9	75.6	0.0
	BS-2012	93.3	100.0	100.0	100.0	33.3	97.8	98.9	98.9	72.2	0.0
	BS-2013	91.1	100.0	100.0	100.0	98.9	97.8	98.9	98.9	98.9	33.3

Table 3: Optimality rates for fractional point guards.

4.3.2 Fractional Guards

The Braunschweig line of work started with solving $\mathrm{AGP}_{\mathbb{R}}(P, P)$, the fractional point guard variant, and all Braunschweig versions, even those designed for binary solutions, still support the fractional AGP. Table 3 shows how often the three implementations could find optimal solutions, and how often they achieved a 5% gap by the end of the 20-minute runtime limit. Here again, the polygons have been grouped: those with holes and those without holes.

Unsurprisingly, there is no significant difference between the `BS-2010` and `BS-2012` versions, the development between these snapshots focused on the integer case. The improvements from `BS-2012` to `BS-2013` stem from improved geometry subroutines which are beneficial to both, the binary and the fractional mode. It can be seen that near-optimal solutions are obtained almost every time, but the gap is not always closed. Furthermore, with the 20-minute time limit, there is an barrier between 2000 and 5000 vertices, where the success rate drops sharply, indicating that the current frontier for input complexity lies roughly in this range.

4.3.3 Point Guards

We turn our attention to the classic AGP, $\mathrm{AGP}(P, P)$: Finding integer solutions with point guards. We report optimality in three different ways: Which percentage of instances could be solved optimally with a matching lower bound (i.e., proven optimality) is reported in Table 4; we show in how many percent of the cases an instance could be solved optimally, whether or not a matching bound was found in Table 5; Table 6 reports how many percent of the solutions were no more than 5% away from the optimum. This allows to distinguish between cases where `BS-2013` does not converge, and cases where the integrality gap prevents it from detecting optimality.

The `C+BS-2013` implementation solves the vast majority of instances from our test set to proven optimality, the only notable exception being some classes of very large polygons with holes and the 5000-vertex Koch polygons. Given how the best known implementation by 2011, the Torino one from Section 3.6, had an optimality rate of about 70% for 60-vertex instances, it is clearly visible how the developments in the last years pushed the frontier. With `C+BS-2013`, instances with 2000 vertices are usually solved to optimality, showing an increase in about two orders of magnitude. The success of `C+BS-2013` is multifactorial: It contains improved combinatorial algorithms as well as faster geometry routines, most notably a fast visibility implementation. Section 5 discusses its key success factors.

It can be seen from Table 6 that many algorithms are able to find near-optimal solutions (5% gap) for most instances, indicating that for practical purposes, all 2013 algorithms perform very well. The frontier on how large instances can be solved with small gap is between 2000 and 5000 vertices for most polygons with holes and beyond 5000 vertices for simple polygons.

Comparing Tables 4–6, it can be seen that the primal-dual approach (`BS-2010` and `BS-2012`) produces decent upper bounds, often optimal ones, but does have an issue with finding matching lower bounds. This drawback has been much improved in `BS-2012` but is still measurable.

Class	n	Optimality Rate (%)				
		BS-2012	C-2013.1	C-2013.2	BS-2013	C+BS-2013
Simple	200	**100.0**	**100.0**	**100.0**	96.7	**100.0**
	500	76.7	**100.0**	**100.0**	96.7	**100.0**
	1000	70.0	96.7	**100.0**	90.0	**100.0**
	2000	36.7	6.7	50.0	60.0	**100.0**
	5000	0.0	0.0	0.0	26.7	**100.0**
Orthogonal	200	96.7	**100.0**	**100.0**	96.7	96.7
	500	86.7	**100.0**	96.7	93.3	93.3
	1000	70.0	**100.0**	**100.0**	86.7	**100.0**
	2000	46.7	70.0	90.0	70.0	**100.0**
	5000	0.0	0.0	0.0	40.0	93.3
Simple-simple	200	93.3	–	**100.0**	86.7	**100.0**
	500	76.7	–	83.3	60.0	**100.0**
	1000	3.3	–	0.0	13.3	**100.0**
	2000	0.0	–	0.0	0.0	46.7
	5000	0.0	–	0.0	0.0	0.0
Ortho-ortho	200	83.3	–	96.7	86.7	**100.0**
	500	53.3	–	83.3	53.3	**100.0**
	1000	16.7	–	3.3	16.7	96.7
	2000	0.0	–	0.0	0.0	33.3
	5000	0.0	–	0.0	0.0	0.0
von Koch	200	**100.0**	**100.0**	**100.0**	**100.0**	**100.0**
	500	**100.0**	96.7	**100.0**	93.3	**100.0**
	1000	**100.0**	46.7	**100.0**	96.7	**100.0**
	2000	83.3	0.0	0.0	86.7	**100.0**
	5000	0.0	0.0	0.0	0.0	0.0
Spike	200	**100.0**	–	**100.0**	96.7	**100.0**
	500	**100.0**	–	**100.0**	**100.0**	**100.0**
	1000	3.3	–	96.7	**100.0**	**100.0**
	2000	0.0	–	96.7	**100.0**	**100.0**
	5000	0.0	–	0.0	96.7	**100.0**

Table 4: Optimality Rate for point guards.

Finally, we analyze how difficult the individual instance classes are. In Tables 4–6, we group them by size and based on whether they feature holes. Table 7 shows optimality rates for each class. We restrict presentation to `BS-2013` here, for the simple reason that it has the highest variation in reported rates.

In each class, we see a continuous decline with increasing input complexity, indicating that local features of an instance play a major role in how hard it is to solve it, rather than this being an intrinsic property of the generator. The only generator that produces "easier" instances than the others is Spike. These are instances tailored for showing the difference between vertex and point guards, requiring few guards to be placed in the middle of certain free areas. We include the Spike instances in our experiments because they are an established class of test cases, being aware that all of the current implementations are able to identify good non-vertex positions for guards, and that this class has to be considered easy.

5 Success Factors

As seen in Section 3, the most effective algorithms for the AGP can be decomposed into four elements:

- Geometric subroutines dealing with computing visibility relations, determining feasibility,

Class	n	Optimality Rate (%) without proof				
		BS-2012	C-2013.1	C-2013.2	BS-2013	C+BS-2013
Simple	200	**100.0**	**100.0**	**100.0**	96.7	**100.0**
	500	80.0	**100.0**	**100.0**	**100.0**	**100.0**
	1000	73.3	**100.0**	**100.0**	**100.0**	**100.0**
	2000	50.0	50.0	80.0	93.3	**100.0**
	5000	0.0	0.0	0.0	83.3	**100.0**
Orthogonal	200	96.7	**100.0**	**100.0**	96.7	96.7
	500	86.7	**100.0**	**100.0**	93.3	93.3
	1000	70.0	**100.0**	**100.0**	90.0	**100.0**
	2000	50.0	96.7	93.3	90.0	**100.0**
	5000	0.0	0.0	0.0	50.0	93.3
Simple-simple	200	96.7	–	**100.0**	90.0	**100.0**
	500	93.3	–	96.7	80.0	**100.0**
	1000	33.3	–	20.0	73.3	**100.0**
	2000	0.0	–	0.0	33.3	50.0
	5000	0.0	–	0.0	0.0	0.0
Ortho-ortho	200	93.3	–	**100.0**	**100.0**	**100.0**
	500	80.0	–	93.3	90.0	**100.0**
	1000	70.0	–	30.0	70.0	96.7
	2000	0.0	–	0.0	30.0	43.3
	5000	0.0	–	0.0	0.0	0.0
von Koch	200	**100.0**	**100.0**	**100.0**	**100.0**	**100.0**
	500	**100.0**	**100.0**	**100.0**	93.3	**100.0**
	1000	**100.0**	70.0	**100.0**	96.7	**100.0**
	2000	83.3	0.0	30.0	90.0	**100.0**
	5000	0.0	0.0	0.0	0.0	0.0
Spike	200	**100.0**	–	**100.0**	96.7	**100.0**
	500	**100.0**	–	**100.0**	**100.0**	**100.0**
	1000	3.3	–	**100.0**	**100.0**	**100.0**
	2000	0.0	–	96.7	**100.0**	**100.0**
	5000	0.0	–	0.0	96.7	**100.0**

Table 5: Optimality Rate without proof for point guards.

- Set Cover subroutines computing (near-)optimal solutions for finite cases,
- Routines to find candidates for discrete guard and witness locations, and
- An outer algorithm combining the three parts above.

In this section, we focus on these techniques.

5.1 Geometric Subroutines

Both groups use the 2D Arrangements package [62] of CGAL which follows the *generic programming paradigm* [5]. For instance, in the case of arrangements it is possible to change the curve type that is used to represent the planar subdivisions or the kernel that provides the essential geometric operations and also determines the number type used. In the context of this work, it is clear that the used curves are simply segments[2]. However, the choice of the geometric kernel can have a significant impact on the runtime.

First of all, it should be noted that among the different kernels that CGAL offers only kernels that provide exact constructions should be considered as any inexact construction is likely to induce inconsistencies in the data structure of the arrangements package. This already

[2]In the context of fading [43] circular arcs may also be required.

Class	n	5% gap Rate in (%)				
		BS-2012	C-2013.1	C-2013.2	BS-2013	C+BS-2013
Simple	200	100.0	100.0	100.0	100.0	100.0
	500	100.0	100.0	100.0	100.0	100.0
	1000	100.0	100.0	100.0	100.0	100.0
	2000	100.0	100.0	96.7	100.0	100.0
	5000	0.0	0.0	0.0	100.0	100.0
Orthogonal	200	100.0	100.0	100.0	100.0	100.0
	500	100.0	100.0	100.0	100.0	100.0
	1000	100.0	100.0	100.0	100.0	100.0
	2000	100.0	100.0	100.0	100.0	100.0
	5000	0.0	0.0	0.0	100.0	100.0
Simple-simple	200	100.0	–	100.0	100.0	100.0
	500	100.0	–	93.3	100.0	100.0
	1000	100.0	–	33.3	100.0	100.0
	2000	0.0	–	0.0	96.7	80.0
	5000	0.0	–	0.0	0.0	0.0
Ortho-ortho	200	100.0	–	100.0	100.0	100.0
	500	100.0	–	100.0	100.0	100.0
	1000	100.0	–	40.0	96.7	100.0
	2000	56.7	–	0.0	76.7	86.7
	5000	0.0	–	0.0	0.0	0.0
von Koch	200	100.0	100.0	100.0	100.0	100.0
	500	100.0	100.0	100.0	100.0	100.0
	1000	100.0	73.3	100.0	100.0	100.0
	2000	100.0	0.0	56.7	100.0	100.0
	5000	0.0	0.0	0.0	3.3	0.0
Spike	200	100.0	–	100.0	96.7	100.0
	500	100.0	–	100.0	100.0	100.0
	1000	3.3	–	96.7	100.0	100.0
	2000	0.0	–	96.7	100.0	100.0
	5000	0.0	–	0.0	96.7	100.0

Table 6: Rate of upper bound within 5% distance to lower bound.

	200	500	1000	2000	5000	Avg.
Simple	96.7	96.7	90.0	60.0	26.7	74.0
Orthogonal	96.7	93.3	86.7	70.0	40.0	77.3
simple-simple	86.7	60.0	13.3	0.0	0.0	32.0
ortho-ortho	86.7	53.3	16.7	0.0	0.0	31.3
von Koch	100.0	93.3	96.7	86.7	0.0	75.3
Spike	96.7	100.0	100.0	100.0	96.7	98.7

Table 7: Optimality rates for BS-2013 on different instance classes.

Class Size	200	500	1000	2000	5000
Simple	1.27	1.46	1.55	1.49	1.35
Orthogonal	1.44	1.60	1.66	1.69	1.65
Simple-simple	2.15	1.72	1.44	1.37	-
Ortho-ortho	1.54	1.30	1.21	1.20	-
von Koch	1.02	1.06	1.10	1.16	-
Spike	1.15	1.61	1.76	2.10	2.56

Table 8: The speedup factor of `C+BS-2013` using the Cartesian kernel and the lazy-exact kernel. Similar numbers were obtained for `BS-2012`. The lazy-exact kernel is now the standard configuration in `BS-2013` and `C+BS-2013`.

holds for seemingly simple scenarios as the code of the arrangement package heavily relies on the assumption that constructions are exact.

This essentially leaves two kernels: The Cartesian kernel and the lazy-exact kernel. For both kernels it is possible to exchange the underlying exact rational number type, but `CGAL::Gmpq` [35] is the recommended one[3].

The **Cartesian kernel**, is essentially the naive application of exact rational arithmetic (using the one that it is instantiated with, in this case `CGAL::Gmpq`). Thus, coordinates are represented by a numerator and denominator each being an integer using as many bits as required. This implies that even basic geometric constructions and predicates are not of constant cost, but depend on the bit-size of their input. For instance, the intersection point of two segments is likely to require significantly more bits than the endpoints of the segments. And this is even more relevant in case of cascaded constructions as the bit growth is cumulative. This effect is very relevant in both approaches due to there iterative nature, e.g., when such a point is chosen to be a new guard or witness position.

The **lazy-exact kernel** [54] tries to attenuate all these effects by using exact arithmetic only when necessary. Every arithmetic operation and construction is first carried out using only double interval arithmetic, that is, using directed rounding, an upper and a lower of the exact value is computed. The hope is that for most cases this is already sufficient to give the correct and certified answer, for instance whether a point is above or below a line. However, for the case when this is not sufficient, each constructed object also knows its history, which makes it possible to carry out the exact rational arithmetic as it is done in the Cartesian kernel in order to determine the correct result. This idea is implemented by the lazy kernel not only on the number type level[4], but also for predicates and constructions, which reduces the overhead (memory and time) that is induced by maintaining the history.

By the genericity of CGAL it is possible to easily exchange the used geometric kernel. Table 8 shows the speedup factors by using the Cartesian kernel vs the lazy-exact kernel for the different instances for `C-2013.1`. It should be noted that all Braunschweig and Campinas implementations since 2007 use a complexity reduction step together with the Cartesian kernel: Whenever a point in a face is generated, it is rounded to a nearby point of lower bit complexity. Without this, neither implementation would be able to solve any instance of substantial size. This speedup technique is missing in the variant with the lazy-exact kernel, as it requires to actually compute the point coordinates before rounding, which would defeat the purpose of the kernel. Therefore the table compare the lazy-exact kernel against the Cartesian kernel with explicit complexity reduction.

[3]Other options are, for instance, `leda::rational` [46] or `CORE::BigRat` [39], but, compared to Gmpq, both imply some overhead and are only recommended in case the usage of the more complex number types of these libraries is required.

[4]This can be achieved by the instantiation of the Cartesian kernel with `CGAL::Lazy_exact_nt<CGAL::Gmpq>`.

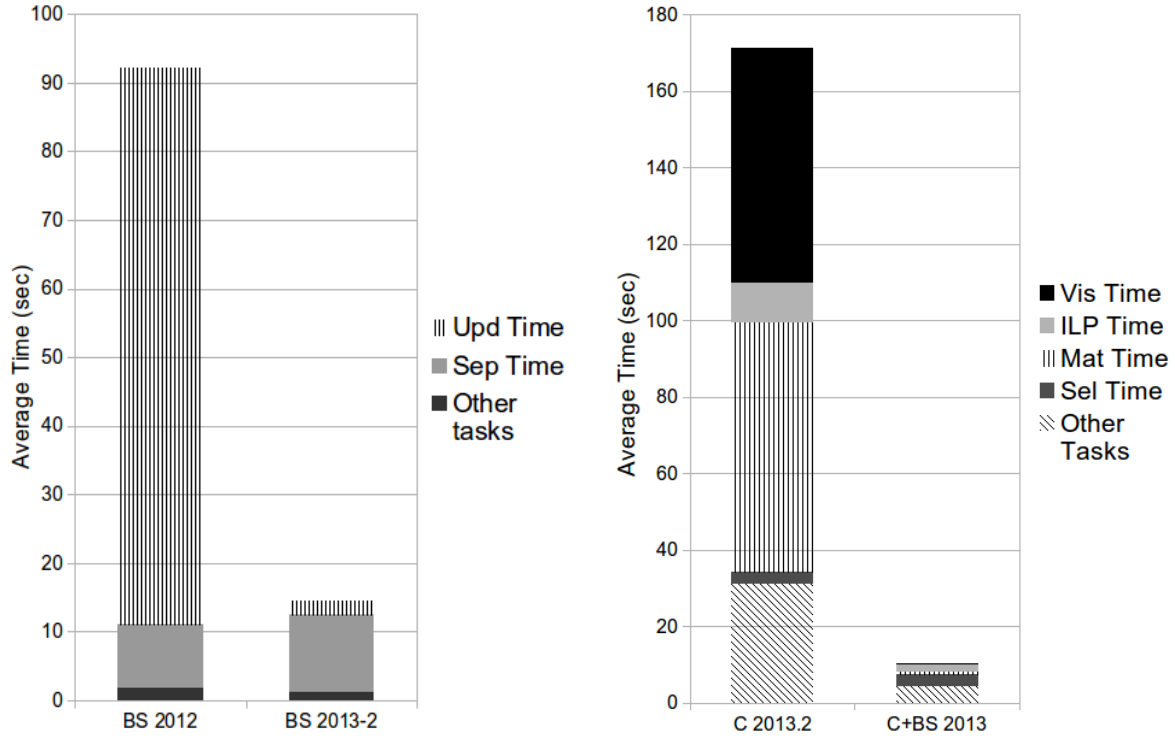

(a) Braunschweig

(b) Campinas

Figure 4: Split up of average total time for different configurations on all simple instances with 1000 vertices. (left) The update time which is dominated by the visibility polygon computation almost vanishes in `BS-2013` compared to the `BS-2012`. (right) The time spent on visibility in `C+BS-2013` is almost negligible compared to the time spend in `C-2013.2`.

For the random polygons, as well as for the spike ones, it can be observed that the lazy-exact kernel is usually almost twice as fast as the Cartesian kernel. However, for the von Koch polygons the lazy-exact kernel only gives a mild speedup. We explain this by two effects. First, the bit-size of the input polygons is not very large and also the bit-size of intermediate constructions do not grow as much, as the horizontal and vertical lines dominate the scene. Second, the instance induces degenerate situations in which the lazy-exact kernel must fall back to the exact arithmetic in which cases effort for interval arithmetic and maintaining the history is a real overhead. The lazy-exact kernel is now the standard configuration in `BS-2013` and `C+BS-2013`.

5.1.1 Visibility Computation

One of the most significant improvements with respect to speed is due to the new upcoming visibility package [36] of CGAL. This package was developed by the group in Braunschweig with this project being the main motivation. Of course, this packages was also made available to the group in Campinas prior to its actual integration in CGAL. Figure 4 illustrates the tremendous impact on the runtime for both approaches. The left side shows the split up of total runtime for the code from Braunschweig in 2012 and 2013. While in 2012 the update time (dominated by visibility computation) used about two third of the time for visibility computation is now almost negligible. The same holds for improvements achieved in the code from Campinas, see right side of Figure 4. It can be noticed in the latter graph that the time spent by `C+BS-2013`

in building the constraint matrices for the ILPs, denoted by `Mat Time`, also suffered a huge reduction relative to `C-2013.2`. As commented in Section 3.11, this was mostly due to the execution of the visibility testing from the perspective of the witnesses rather than the guards.

5.2 Set Cover Optimization

Many AGP algorithms rely on repeatedly solving AGP(G, W) (Equations (1)–(3)) for finite G and W as a subroutine, corresponding to the NP-hard SCP. Therefore improving the solutions times for these SCP instances can benefit the overall algorithm.

5.2.1 Lagrangian Relaxation

In the algorithm developed by the research group in Campinas subsequent to the journal version from 2013 [59] (Section 3.9), attempts were made to reduce the time spent by the ILP solver through the implementation of some known techniques, such as ILP matrix reduction and Lagrangian heuristic.

A standard method for reducing constraints and variables was used, which is based on inclusion properties among columns (guard candidates) and rows (witnesses) of the Boolean constraint matrix of the ILP that models the SCP instance.

Furthermore, their algorithm employs a *Lagrangian Heuristic* in order to obtain good, hopefully optimal, feasible starting solutions for the SCP to speedup the convergence towards an optimum. See [8] for a comprehensive introduction to this technique. The heuristic implemented is based on the work presented in [8]. Figure 5 shows how the use of this technique positively influenced the average run time of the approach.

5.2.2 DC Programming

A different solution method for the Braunschweig approach was discussed in Kröller et al. [42]. Here, the ILP representing the SCP for AGP(G, W) was rewritten as

$$\min_{x \in \mathbb{R}^G} F(x) \text{ , where } F(x) := \sum_{g \in G} x_g - \theta \sum_{g \in G} x_g(x_g - 1) + \chi(x) \text{ .} \tag{5}$$

Here, θ is a sufficiently large constant used to penalize fractional values for x_g, and $\chi \colon \mathbb{R}^G \to \{0, \infty\}$ is an indicator function with

$$\chi(x) = 0 \ :\Longleftrightarrow \ \begin{cases} \sum_{g \in \mathcal{V}(w)} x_g \geq 1 & \forall w \in W \\ 0 \leq x_g \leq 1 & \forall g \in G \end{cases} \text{ .} \tag{6}$$

It is easy to see that F can be expressed as $F(x) := f_1(x) - f_2(x)$, where

$$f_1(x) := \sum_{g \in G} x_G + \chi(x), \quad \text{and} \quad f_2(x) = \theta \sum_{g \in G} x_g(x_g - 1) \text{ ,} \tag{7}$$

i.e., the SCP instance is reduced to minimizing the difference of two non-linear convex functions. For such optimization problems, the DCA algorithm [53] can be used. In experiments, it was shown that solutions for AGP(G, W) could be found very quickly, however, at the time, the large runtime overhead of the geometric subroutines led to inconclusive results on the benefits. Revisiting this approach with the new `BS-2013` and `C+BS-2013` implementations, which no longer suffer from this overhead, will be an interesting experiment left for future work.

5.3 Point Generation

The central heuristic component in most AGP implementations are point generators, which choose where to place new guards and witnesses. One cannot expect these problems to be simple, given that a perfect choice for G and W equals solving AGP optimally.

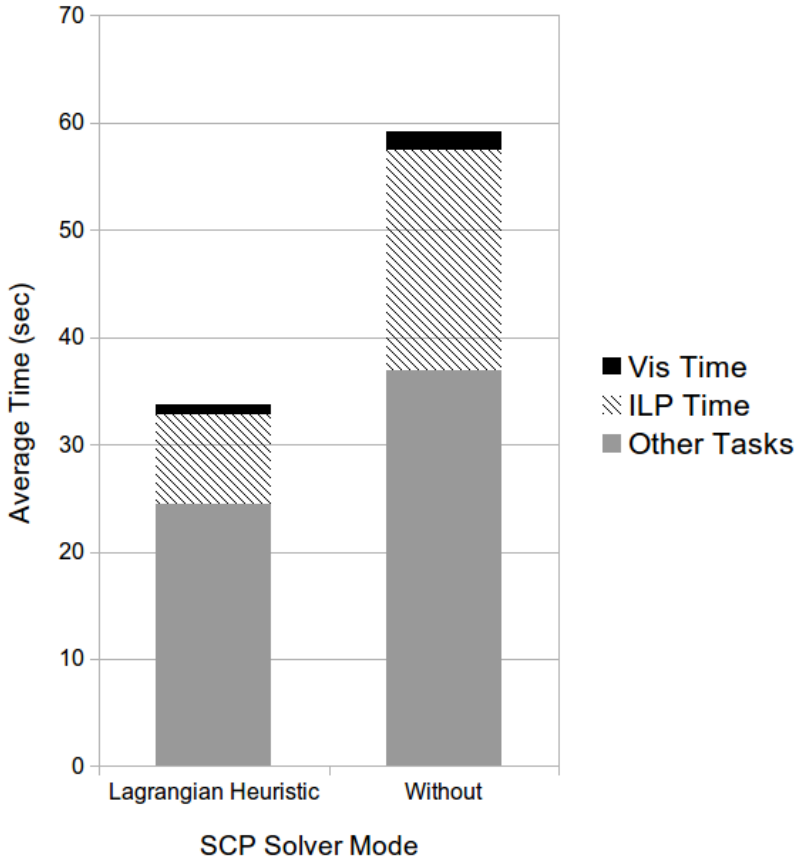

Figure 5: Average Time needed for the current Campinas version to solve von Koch polygons with 1000 vertices with and without the Lagrangian Heuristic.

5.3.1 Guard Placement

One subroutine in the algorithms is to improve the current set of guards, given a current set W of witnesses. This corresponds to finding guards that can be used to solve $\mathrm{AGP}(P, W)$.

A critical observation [60] allows for elegant solution to this problem: Consider the visibility arrangement $\mathcal{A}(W)$. It is always possible to find an optimal solution for $\mathrm{AGP}(P, W)$ where each AVP contains at most one guard. This can be strengthened by observing that the guards can be restricted further to light AVPs. As explained in Section 3.8, the `C-2013.1` algorithm uses as guard candidates the vertices of P along with all vertices from light AVPs. In `C+BS-2013`, a second guard placement strategy using no more than one interior point per AVP is available. Results comparing these two can be seen in Table 9. It is possible to conclude that the latest guard placement strategy, which consists of using only one point within each light AVP, is often the best option. The explanation for this success is probably related to the fact that, with the winning strategy, there is a reduced number of visibility tests between witnesses and guard candidates, as well as a smaller size of SCP instances to be solved.

For $\mathrm{AGP}_{\mathbb{R}}(P, W)$, as solved by the Braunschweig line of algorithms, this observation can be extended further: If an optimal dual solution for $\mathrm{AGP}_{\mathbb{R}}(G, W)$ is available, selecting additional guards corresponds to a column generation process. Therefore, the BS algorithms place guards only in light AVPs where the dual solution guarantees an improvement in the objective function. To avoid cycling in the column generation process, G is monotonically growing, leading over time to a large number of guard positions.

	Vertices of light AVPs	Interior of light AVPs
simple 2000	100.0	100.0
ortho 2000	100.0	100.0
simple-simple 2000	6.7	33.3
ortho-ortho 2000	13.3	46.7
von Koch 2000	100.0	100.0
spike 2000	100.0	100.0

Table 9: Percentage of instances solved to binary optimality by the current implementation from Campinas with guard candidates on vertices or inside light AVPs.

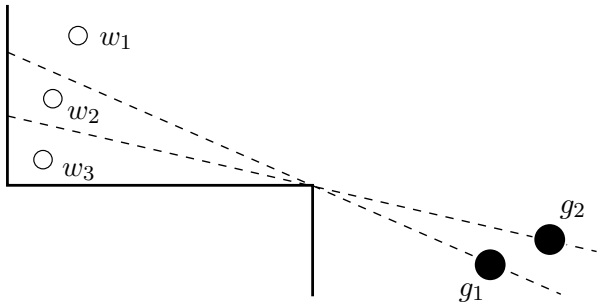

Figure 6: Creeping shadow effect.

5.3.2 Witness Placement

The choice of witnesses is as important as that of the guards. In principle, the same reasoning as for guards can be used: Given guard candidates G, creating W with one witness in every shadow AVP of $\mathcal{A}(G)$ guarantees that a solution for $\mathrm{AGP}(G, W)$ is also a solution for $\mathrm{AGP}(G, P)$. A naïve placement algorithm based on this observation would simply create witnesses in shadow AVPs. However, this leads to the problem of creeping shadows at reflex vertices, see Figure 6: Placing witness in the interior of the AVP adjacent to the polygon boundary creates an infinite chain of guard/witness positions that converges towards a witness on the boundary, but not reaching it. Both the Braunschweig and the Campinas algorithms therefore can create additional witnesses on the edges of shadow AVPs.

5.3.3 Initial Set

The selection of the first candidates for guards and witnesses, i.e., the initial choice of G and W can have tremendous impact on algorithm runtime. In principle, a good heuristic here could pick an almost optimal set for G and a matching W to prove it, and reduce the algorithm afterwards to a few or even no iterations.

Chwa et al. [14] provide a partial answer to this problem: They attempt to find a finite set of witnesses with the property that guarding this set guarantees guarding the whole polygon. If such a set exists, the polygon is called *witnessable*. Unfortunately this is not always the case. However, for a witnessable polygon, the set can be characterized and quickly computed. Current algorithms do not bother checking for witnessability (although Chwa et al. provide an algorithm), but rather directly compute this set and use it for initial witnesses. Should the polygon be witnessable, the algorithm automatically terminates in the first iteration.

Considering the current version from Campinas, two initial discretizations are used: Convex Vertices (CV) and Chwa Points (CP). The first one includes only the convex vertices of the polygon in the initial set, while the second chooses the middle points of reflex-reflex edges and the convex vertices that are adjacent to reflex vertices.

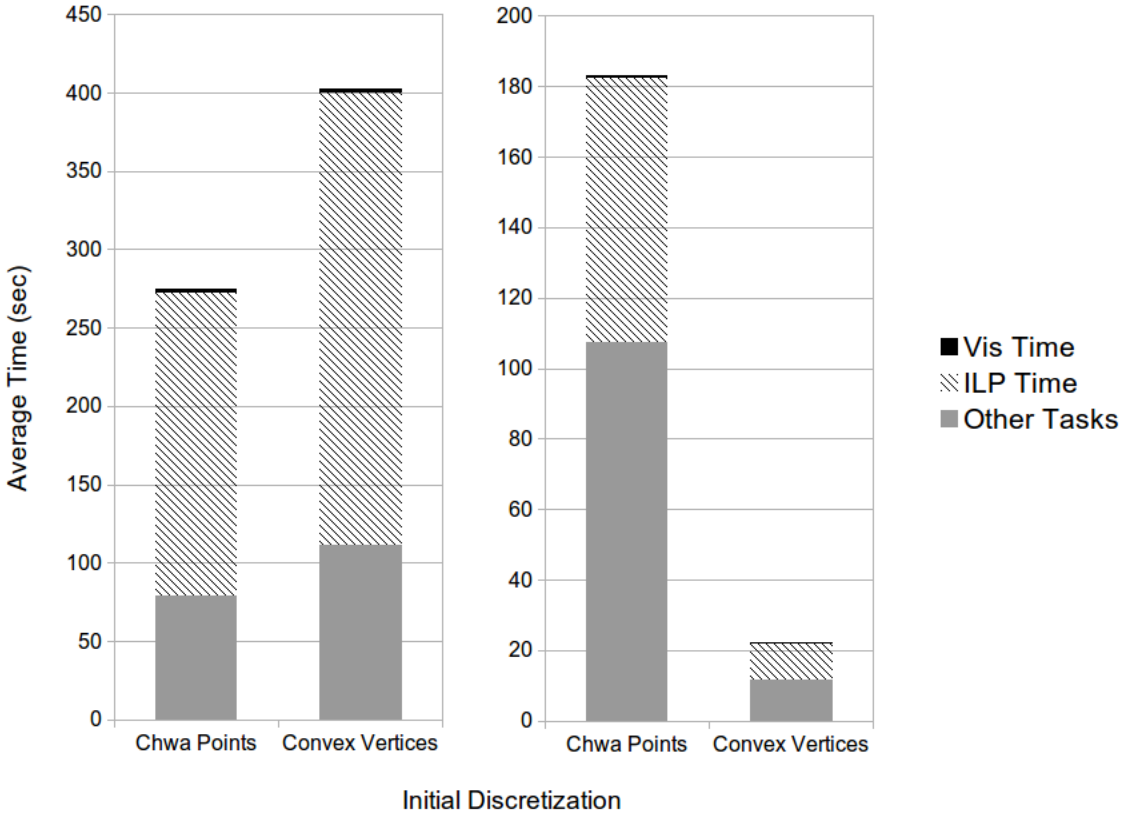

Figure 7: Average Time needed to solve ortho-ortho (left) and spike (right) polygons with 1000 vertices using the Convex Vertices and the Chwa Points discretization.

The two charts in Figure 7 show the average run time necessary to find optimal solutions when using CV and CP strategies on simple-simple and spike polygons. From these charts, one can perceive that there is an advantage in using the CP discretization for polygons from the simple-simple class. On the other hand, the chart corresponding to the spike polygons shows that the implementation works much better when the strategy chosen is the CV one. In this last case, the program required four times less time to solve the same set of polygons when using the CV strategy as opposed to CP.

For `BS-2010`, several strategies were implemented, see Table 10 which is extracted from the corresponding paper [41]: Leaving G and W empty (for implementation reasons, both contained one arbitrary point), putting guards and witnesses on every (or every other) vertex of the polygon, putting guards on all reflex vertices, and putting a witness on every edge adjacent to a reflex vertex. The Chwa-inspired combination allowed for a speedup of around two.

5.4 Lower Bounds

A crucial success factor for solving the binary AGP variants is the quality of the lower bounds. This is especially visible in `BS-2013`, which was tested with and without the cutting planes, i.e., with and without the features published in [33, 31] and outlined in Section 3.7. Table 11 compares the solution rates for the different classes of instances with 500 vertices and clearly shows that using cutting planes greatly improves solution rates. Cutting planes increase the lower bounds and improve the solution rates for all classes of instances.

For the Campinas approach, the quality of the lower bound computed is a very important

Initial G	Initial W	Speedup
Single Point	Single Point	1.00
Every other vertex	Every other vertex	1.59
All vertices	All vertices	1.64
All vertices	Reflex edges	1.74
Reflex vertices	Reflex edges	2.02

Table 10: Speedup factors in `BS-2010` obtained by varying initial guards and witnesses [41].

Class / Technique	With Cutting Planes	Without Cutting Planes
ortho	80.0%	63.3%
simple	86.7%	40.0%
von Koch	100.0%	70.0%
ortho-ortho	63.3%	13.3%
simple-simple	70.0%	6.7%
spike	100.0%	96.7%

Table 11: Percentage of instances solved to binary optimality comparing two variants of code from Braunschweig 2013, one with and without cutting planes, for 500-vertex instances.

issue. For AGP(P, P), the lower bound is obtained by solving an AGP(P, W) instance, where W is a discretized set of witnesses points within P. Therefore, it is fair to say that the quality of the value computed is directly dependent on the strategy applied to select the points that comprise the set W. For more information on how the witness set is managed and how it affects convergence of Campinas method, see Section 5.3.

6 Variants and Open Problems

6.1 Fading

An interesting variant for the AGP was proposed by Joe O'Rourke in 2005: What if visibility suffers from fading effects, just like light in the real world does? To be precise, we assume that for a guard g with intensity x_g, a witness $w \in \mathcal{V}(g)$ is illuminated with a value of $\varrho(d(g, w))x_g$, where $d(g, w)$ is the Euclidean distance between g and w, and ϱ is a fading function, usually assumed to be

$$\varrho(d) := \begin{cases} 1 & \text{if } d < 1 \\ d^{-\alpha} & \text{if } 1 \leq d < R \\ 0 & \text{if } d \geq R \end{cases} . \tag{8}$$

Here, α is a constant (2 for natural light in 3D space), and R is a maximal radius beyond which illumination is neglected. Fixing $\varrho(d)$ to 1 for small d is necessary to keep the problem well-defined. Otherwise, an infinitesimally small light can illuminate a small circle around it. Then, no finite solution can exist, because it can always be improved by creating additional guards between the existing ones, and reducing intensity for all. This converges towards the setup of $G = P$, with all $x_g = 0$, which is not feasible.

Very little is known about this variant. A restricted case has been discussed by Eisenbrand et al. [26], where a 1-dimensional line segment is illuminated from a fixed set of guards. It is shown how to solve this problem exactly and approximatively using techniques from mathematical programming.

The primal-dual Braunschweig algorithm was shown to apply to this variant as well: Kröller

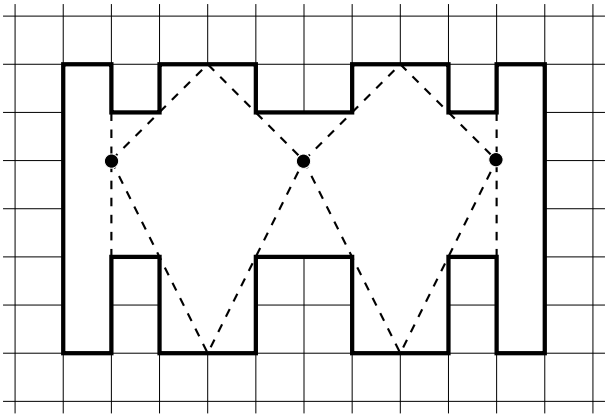

Figure 8: A simple orthogonal polygon possessing only a single optimal solution.

et al. [43] have modified the ILP formulation (1)–(3) to use the constraint

$$\sum_{g \in \mathcal{V}(w)} \varrho(d(g,w)) x_g \geq 1 \quad \forall w \in W \tag{9}$$

instead of (2). Two algorithms for vertex guards were proposed and tested [29], based on the `BS-2013` implementation. The first approximates ϱ with a step function, and uses updated primal and dual separation routines that operate on overlays of visibility polygons and circular arcs, resulting in an FPTAS for the fractional AGP(V, P). The other is based on continuous optimization techniques, namely a simplex partitioning approach. In an experimental evaluation using polygons with up to 700 vertices, it was found that most polygons can be solved (to an 1.2-approximation in case of the discrete approach) within 20 minutes on a standard PC. The continuous algorithm turned out to be much faster, and very often finishing with an almost-optimal solution with a gap under 0.01%. In an experimental work by Kokemüller [40], AGP(P, P) with fading was analyzed. It was found that placing guards makes the problem substantially more difficult. This is mainly due to an effect where moving one guard requires moving chains of other guards as well to cover up for decreased illumination. It was also found that scaling an input polygon has an impact on the structure of solutions and number of required guards, resulting in a dramatic runtime impact.

6.2 Degeneracies

The experiments conducted by different groups as well as the results shown in Section 4 indicate that practically efficient algorithms exist, and a growing number of input instances can be solved to optimality. This raises the question whether it can be expected that all instances can be solved, given sufficient time.

Unfortunately, the answer to this question is "no". As a counterexample, consider the polygon depicted in Figure 8. The three indicated guard positions form the only optimal solution. There is no variation allowed—shifting any guard by any $\varepsilon > 0$, in an arbitrary direction, will create a shadow, requiring a fourth guard and thereby losing optimality.

None of the currently known algorithms can solve such problems, as no way to characterize these points is known. To see this, consider perturbations of the shown polygon: It is possible to slightly move all vertices in a way that keeps the dashed lines intact. It is not clear how to find the shadow alignment points on the boundary, which in turn define the optimal guard positions. It should be noted, however, that it remains an open question whether there are polygons given by rational coordinates that require optimal guard positions with irrational coordinates.

To summarize, after forty years of research on AGP, it is still not known whether there exist finite-time algorithms for it. Even membership in NP is unclear, as it is not known if guard

locations can be encoded in polynomial size.

7 Conclusion

In this paper, we have surveyed recent developments on solving the Art Gallery Problem (AGP) in a practically efficient manner. After over thirty years of mostly theoretical work, several approaches have been proposed and evaluated over the last few years, resulting in dramatic improvements. The size of instances for which optimal solutions can be found in reasonable time has improved from tens to thousands of vertices in just a few years.

While these developments are very promising, experimental findings have led to new questions about the problem complexity. There are bad instances that current implementations cannot solve despite small size, and it is not clear whether exact algorithms for the AGP can exist, even ones with exponential runtime.

Acknowledgments

Many people have contributed to the developments described in this paper. In particular, the authors would like to thank Tobias Baumgartner, Marcelo C. Couto, Sándor P. Fekete, Winfried Hellmann, Mahdi Moeini, Eli Packer, and Christiane Schmidt.

Stephan Friedrichs was affiliated with TU Braunschweig, IBR during most of the research.

This work was partially supported by the Deutsche Forschungsgemeinschaft (DFG) under contract number KR 3133/1-1 (Kunst!), by Fundação de Amparo à Pesquisa do Estado de São Paulo (Fapesp, #2007/52015-0, #2012/18384-7), Conselho Nacional de Desenvolvimento Científico e Tecnológico (CNPq, grants #311140/2014-9, #477692/2012-5 and #302804/2010-2), and Faepex/Unicamp. Google Inc. supported the development of the Computational Geometry Algorithms Library [12] (CGAL) visibility package through the 2013 Google Summer of Code.

References

[1] M. Aigner and G. M. Ziegler. *Proofs from THE BOOK*. Springer Publishing Company, Incorporated, 4th edition, 2009.

[2] Y. Amit, J. S. B. Mitchell, and E. Packer. Locating guards for visibility coverage of polygons. In *ALENEX*, pages 120–134, 2007.

[3] Y. Amit, J. S. B. Mitchell, and E. Packer. Locating guards for visibility coverage of polygons. *International Journal of Computational Geometry & Applications*, 20(5):601–630, 2010.

[4] T. Asano. An efficient algorithm for finding the visibility polygon for a polygonal region with holes. *IEICE Transactions*, 68(9):557–559, 1985.

[5] M. H. Austern. *Generic Programming and the STL*. PUB-AW, 1999.

[6] E. Balas and S. M. Ng. On the set covering polytope: II. lifting the facets with coefficients in {0, 1, 2}. *Mathematical Programming*, 45:1–20, 1989. 10.1007/BF01589093.

[7] T. Baumgartner, S. P. Fekete, A. Kröller, and C. Schmidt. Exact solutions and bounds for general art gallery problems. In *Proceedings of the SIAM-ACM Workshop on Algorithm Engineering and Experiments, ALENEX 2010*, pages 11–22. SIAM, 2010.

[8] J. E. Beasley. Lagrangian relaxation. In C. R. Reeves, editor, *Modern heuristic techniques for combinatorial problems*, pages 243–303. John Wiley & Sons, Inc., New York, NY, USA, 1993.

[9] A. Bottino and A. Laurentini. A nearly optimal sensor placement algorithm for boundary coverage. *Pattern Recognition*, 41(11):3343–3355, 2008.

[10] A. Bottino and A. Laurentini. A nearly optimal algorithm for covering the interior of an art gallery. *Pattern Recognition*, 44(5):1048–1056, 2011.

[11] F. Bungiu, M. Hemmer, J. Hershberger, K. Huang, and A. Kröller. Efficient computation of visibility polygons. *CoRR*, abs/1403.3905, 2014.

[12] CGAL (Computational Geometry Algorithms Library). `http://www.cgal.org/`.

[13] V. Chvátal. A combinatorial theorem in plane geometry. *Journal of Combinatorial Theory Series B*, 18:39–41, 1974.

[14] K. Chwa, B. Jo, C. Knauer, E. Moet, R. van Oostrum, and C. Shin. Guarding art galleries by guarding witnesses. *International Journal of Computational Geometry & Applications*, 16(2-3):205–226, 2006.

[15] M. C. Couto, P. J. de Rezende, and C. C. de Souza. An exact algorithm for an art gallery problem. Technical Report IC-09-46, Institute of Computing, University of Campinas, Nov. 2009.

[16] M. C. Couto, P. J. de Rezende, and C. C. de Souza. Instances for the Art Gallery Problem, 2009. http://www.ic.unicamp.br/∼cid/Problem-instances/Art-Gallery.

[17] M. C. Couto, P. J. de Rezende, and C. C. de Souza. An IP solution to the art gallery problem. In *SoCG '09: Proceedings of the 25th Annual Symposium on Computational geometry*, pages 88–89, New York, NY, USA, 2009. ACM.

[18] M. C. Couto, P. J. de Rezende, and C. C. de Souza. Video: An IP solution to the art gallery problem. 18th Video Review of Computational Geometry at the 25th Annual Symposium on Computational Geometry, June 2009. www.computational-geometry.org/SoCG-videos/socg09video/video1-couto.mov.

[19] M. C. Couto, P. J. de Rezende, and C. C. de Souza. An exact algorithm for minimizing vertex guards on art galleries. *International Transactions in Operational Research*, 18:425–448, 2011.

[20] M. C. Couto, C. C. de Souza, and P. J. de Rezende. An exact and efficient algorithm for the orthogonal art gallery problem. In *SIBGRAPI '07: Proceedings of the XX Brazilian Symposium on Computer Graphics and Image Processing*, pages 87–94, Washington, DC, USA, 2007. IEEE Computer Society.

[21] M. C. Couto, C. C. de Souza, and P. J. de Rezende. Experimental evaluation of an exact algorithm for the orthogonal art gallery problem. In C. C. McGeoch, editor, *Proceedings of the 7th International Workshop on Experimental Algorithms (WEA '08)*, volume 5038 of *Lecture Notes in Computer Science*, pages 101–113. 2008.

[22] IBM ILOG CPLEX Optimization Studio. `http://www.ibm.com/software/integration/optimization/cplex-optimizer/`.

[23] A. Deshpande, T. Kim, E. D. Demaine, and S. E. Sarma. A pseudopolynomial time o(logn)-approximation algorithm for art gallery problems. In F. Dehne, J.-R. Sack, and N. Zeh, editors, *Algorithms and Data Structures*, volume 4619 of *Lecture Notes in Computer Science*, pages 163–174. Springer Berlin Heidelberg, 2007.

[24] A. Efrat and S. Har-Peled. Guarding galleries and terrains. *Information Processing Letters*, 100(6):238–245, 2006.

[25] S. Eidenbenz, C. Stamm, and P. Widmayer. Inapproximability results for guarding polygons and terrains. *Algorithmica*, 31(1):79–113, 2001.

[26] F. Eisenbrand, S. Funke, A. Karrenbauer, and D. Matijevic. Energy-aware stage illumination. *Proceedings of the 21st ACM Symposium on Computational Geometry (SCG '05)*, pages 336–345, 2005.

[27] U. M. Erdem and S. Sclaroff. Optimal placement of cameras in floorplans to satisfy task requirements and cost constraints. In *Proceedings of the Fifth International Workshop on Omnidirectional Vision, Camera Networks and Non-classical Cameras*, pages 30–41, 2004.

[28] U. M. Erdem and S. Sclaroff. Automated camera layout to satisfy task-specific and floor plan-specific coverage requirements. *Computer Vision and Image Understanding*, 103(3):156–169, 2006.

[29] M. Ernestus, S. Friedrichs, M. Hemmer, J. Kokemüller, A. Kröller, M. Moeini, and C. Schmidt. Algorithms for Art Gallery Illumination. *ArXiv e-prints*, Oct. 2014.

[30] S. P. Fekete, S. Friedrichs, A. Kröller, and C. Schmidt. Facets for art gallery problems. In D.-Z. Du and G. Zhang, editors, *Computing and Combinatorics*, volume 7936 of *Lecture Notes in Computer Science*, pages 208–220. Springer Berlin Heidelberg, June 2013.

[31] S. P. Fekete, S. Friedrichs, A. Kröller, and C. Schmidt. Facets for art gallery problems. *Algorithmica*, 73(2):411–440, 2014.

[32] S. Fisk. A short proof of Chvátal's watchman theorem. *Journal of Combinatorial Theory, Series B*, 24(3):374–375, 1978.

[33] S. Friedrichs. Integer solutions for the art gallery problem using linear programming. Master's thesis, TU Braunschweig, 2012.

[34] S. K. Ghosh. Approximation algorithms for art gallery problems in polygons and terrains. In M. Rahman and S. Fujita, editors, *WALCOM: Algorithms and Computation*, volume 5942 of *Lecture Notes in Computer Science*, pages 21–34. Springer Berlin / Heidelberg, 2010.

[35] GNU Multiple Precision Arithmetic Library, 2013. http://gmplib.org.

[36] M. Hemmer, K. Huang, and F. Bungiu. 2d visibility. In *CGAL User and Reference Manual*. CGAL Editorial Board, to appear edition, 2014.

[37] R. Honsberger. *Mathematical Gems II*. Mathematical Association of America, 1976.

[38] J. Kahn, M. Klawe, and D. Kleitman. Traditional art galleries require fewer watchmen. *SIAM Journal on Algebraic and Discrete Methods*, 4(2):194–206, 1983.

[39] V. Karamcheti, C. Li, I. Pechtchanski, and C. Yap. A core library for robust numeric and geometric computation. In *Proceedings of the 15th Annual ACM Symposium of Computational Geometry (SCG)*, pages 351–359, 1999.

[40] J. Kokemüller. Variants of the art gallery problem. Master's thesis, TU Braunschweig, 2014.

[41] A. Kröller, T. Baumgartner, S. P. Fekete, and C. Schmidt. Exact solutions and bounds for general art gallery problems. *ACM Journal of Experimental Algothmmics*, 17(2.3), 2012.

[42] A. Kröller, M. Moeini, and C. Schmidt. A novel efficient approach for solving the art gallery problem. In *Seventh International Workshop on Algorithms and Computation (WALCOM '13)*, pages 5–16, 2013.

[43] A. Kröller and C. Schmidt. Energy-aware art gallery illumination. In *Proceedings of the 28th European Workshop on Computational Geometry (EuroCG '12)*, pages 93–96, 2012.

[44] A. Laurentini. Guarding the walls of an art gallery. *The Visual Computer*, 15(6):265–278, 1999.

[45] D.-T. Lee and A. K. Lin. Computational complexity of art gallery problems. *IEEE Transactions on Information Theory*, 32(2):276–282, 1986.

[46] K. Mehlhorn and S. Näher. LEDA*: A Platform for Combinatorial and Geometric Computing*. PUB-CAMB, Cambridge, UK, 2000.

[47] J. S. B. Mitchell. Approximating watchman routes. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013*, pages 844–855, 2013.

[48] B. J. Nilsson. Approximate guarding of monotone and rectilinear polygons. In L. Caires, G. Italiano, L. Monteiro, C. Palamidessi, and M. Yung, editors, *Automata, Languages and Programming*, volume 3580 of *Lecture Notes in Computer Science*, pages 1362–1373. Springer Berlin / Heidelberg, 2005.

[49] J. O'Rourke. *Art Gallery Theorems and Algorithms*. International Series of Monographs on Computer Science. Oxford University Press, New York, NY, 1987.

[50] J. O'Rourke and K. Supowit. Some NP-hard polygon decomposition problems. *IEEE Transactions on Information Theory*, 29(2):181–190, Mar. 1983.

[51] E. Packer. Computing multiple watchman routes. In C. C. McGeoch, editor, *Proceedings of the 7th International Workshop on Experimental Algorithms, WEA 2008, Provincetown, MA, USA, May 30-June 1, 2008*, volume 5038 of *Lecture Notes in Computer Science*, pages 114–128. 2008.

[52] E. Packer. *Robust Geometric Computing and Optimal Visibility Coverage*. PhD thesis, SUNY Stony Brook, 2008.

[53] T. Pham Dinh and H. Le Thi. Convex analysis approach to D.C. programming: Theory, algorithms and applications. *Acta Mathematica Vietnamica*, 22(1):289–355, 1997.

[54] S. Pion and A. Fabri. A generic lazy evaluation scheme for exact geometric computations. In *2nd # WOR-LCSD*, 2006.

[55] D. Schuchardt and H.-D. Hecker. Two NP-hard art-gallery problems for ortho-polygons. *Mathematical Logic Quarterly*, 41:261–267, 1995.

[56] T. C. Shermer. Recent results in art galleries (geometry). *Proceedings of the IEEE*, 80(9):1384–1399, Sept. 1992.

[57] A. P. Tomás, A. L. Bajuelos, and F. Marques. Approximation algorithms to minimum vertex cover problems on polygons and terrains. In *Proceedings of the International Conference on Computational Science (ICCS 2003)*, volume 2657 of *Lecture Notes in Computer Science*, pages 869–878. Springer Berlin / Heidelberg, 2003.

[58] A. P. Tomás, A. L. Bajuelos, and F. Marques. On visibility problems in the plane - solving minimum vertex guard problems by successive approximations. In *Proceedings of the 9th International Symposium on Artificial Intelligence and Mathematics (AI & MATH 2006)*, 2006. to appear.

[59] D. C. Tozoni, P. J. de Rezende, and C. C. de Souza. A practical iterative algorithm for the art gallery problem using integer linear programming. *Optimization Online*, October 2013. www.optimization-online.org/DB_HTML/2013/11/4106.html.

[60] D. C. Tozoni, P. J. de Rezende, and C. C. de Souza. The quest for optimal solutions for the art gallery problem: A practical iterative algorithm. In V. Bonifaci, C. Demetrescu, and A. Marchetti-Spaccamela, editors, *Experimental Algorithms*, volume 7933 of *Lecture Notes in Computer Science*, pages 320–336. Springer Berlin Heidelberg, 2013.

[61] J. Urrutia. Art gallery and illumination problems. In J.-R. Sack and J. Urrutia, editors, *Handbook on Computational Geometry*, pages 973–1026. Elsevier Science Publishers, 2000.

[62] R. Wein, E. Berberich, E. Fogel, D. Halperin, M. Hemmer, O. Salzman, and B. Zukerman. 2D arrangements. In *CGAL User and Reference Manual*. CGAL Editorial Board, 4.0 edition, 2012.